

Bridging AI's trust gaps

Aligning policymakers
and companies

July 2020



THE
FUTURE
SOCIETY

Contents

Foreword	1
Background	2
Key survey findings	4
Policymakers have a clear vision of AI ethical risks – and are moving to implementation	5
Companies are misaligned with policymakers at this critical moment	7
Companies’ misalignment with policymakers creates new risks	10
Companies and policymakers have contrasting gaps	12
These contrasting gaps require a comprehensive approach to regulation	13
A trust gap complicates matters – and creates an opening for third parties	14
Implications for companies	16
Implications for policymakers	17
Appendix: Methodology	18

Exhibits

Key survey findings	
Exhibit 1. Facial recognition check-ins: policymakers have a clear vision of AI ethical principles	5
Exhibit 2. Home virtual voice assistants: policymakers have a clear vision of AI ethical principles	6
Exhibit 3. Companies rate “privacy and data rights” the most important issue	7
Exhibit 4. Companies rate “safety and security” more highly than do policymakers	8
Exhibit 5. Policymakers and companies disagree on who will lead a multi-stakeholder framework	9
Exhibit 6. How much variation will likely occur in developing and implementing AI ethical principles across countries/regions?	10
Exhibit 7. Policymakers are less immersed in technical and business challenges	12
Exhibit 8. Companies have a trust gap with policymakers	14
Appendix: Methodology	
Exhibit A1. Use cases	20
Exhibit A2. Ethical principles	21

FOREWORD

The rapid development of artificial intelligence (AI) is raising urgent questions about ethical and consumer protection issues — from potential bias in algorithmic recruiting decisions to the privacy implications of health monitoring applications.

Policymakers have been exploring these issues in depth and are well aligned on identifying the most important ethical principles. They have now reached a pivotal point: moving from articulating principles to implementing public policy.

The COVID-19 global pandemic has highlighted complex ethical trade-offs, reoriented the ethical landscape and inspired unprecedented coordination between policymakers and companies. AI-based solutions are playing a crucial role in the fight against the pandemic — from contact-tracing algorithms to drones that scan body temperatures and ensure compliance with social distancing directives. The primacy of tackling the crisis is reordering ethical considerations.

Concerns about fairness and privacy, which policymakers rated as the top ethical issues in the use of AI for law enforcement surveillance when responding to our survey in December and January, now appear to be downplayed. Instead, in the face of a pervasive lethal challenge, safety and security is being given greater weight. Meanwhile, as scientists and regulators race to find new treatments, the principle of innovation is becoming relatively more important, while the need for precaution is being toned down.

It's possible these softening norms could have some staying power. On the other hand, concerns about privacy could instead be intensified after the crisis if the pandemic makes technologies such as facial recognition or health tracking apps more widespread.

The COVID-19 crisis appears to have reinforced policymakers' momentum to enact policies addressing AI ethics. As such, the findings of this study are timely and urgent. At this critical juncture — when policymakers are moving to implementation and a global pandemic could heighten ethical concerns — companies are misaligned with policymakers on key ethical issues. This lack of understanding creates a host of new risks for companies. Compounding the challenge, companies also have a trust deficit with policymakers, making it difficult to achieve much-needed coordination.

This survey provides relevant insights for leaders in business and government. Closing the gaps identified here should be a top priority. Doing so would help companies mitigate risks, help policymakers enact policies that are nuanced and realistic — and, most important, address users' concerns about the ethical implications of AI.





BACKGROUND

AI promises to transform our lives with applications ranging from autonomous vehicles to virtual assistants. While we are at the dawn of the AI age, its potential applications are already triggering unease in the public and receiving heightened scrutiny from policymakers.

These concerns are emerging against the backdrop of a “techlash” fueled by growing anxiety about privacy breaches, data rights, polarization and disinformation. In recent years, policymakers have been responding to these concerns. Most notably, the European Union passed the General Data Protection Regulation (GDPR) which went into effect in 2018. Policymakers in several other jurisdictions are either examining these issues or have passed similar laws of their own, from Brazil’s Lei Geral de Proteção de Dados (LGPD) to the California Consumer Privacy Act (CCPA).

As AI scales up in new applications, it is fueling similar fears – as well as entirely new ones. Will the algorithms powering autonomous vehicles keep passengers safe? Will automated loan application decisions be transparent and non-discriminatory? Will facial recognition cameras violate citizens’ privacy or mistakenly target innocent individuals?

Policymakers have been responding to these concerns. So far, they have articulated ethical principles to guide the trustworthy development and adoption of AI. Since 2016, more than 100 ethical guidelines have been published by governmental bodies, multi-stakeholder groups, academic institutions and private companies. Now, momentum is growing to put these principles into practice through regulation or other policy means.¹ In February 2020, the European Commission released a White Paper on AI detailing its comprehensive approach on these issues.



However, for these efforts to be fruitful, companies and policymakers need to be aligned. Coordination between both sets of stakeholders is critical to develop pragmatic policy and governance approaches that are informed by constraints and realities on the ground.

If alignment between policymakers and companies is paramount, the inevitable question is: how aligned are they currently? What issues are policymakers prioritizing as they focus on emerging applications of AI? To what extent are companies aware of, and planning for, these new issues?

To answer these questions, we undertook a global survey of policymakers² and companies in late 2019 and early 2020. The survey reveals key gaps in alignment between the private sector and policymakers, which create new market and legal risks. As we move from principles to practice, greater coordination and collaboration is necessary.

The survey covered 12 use cases (applications of AI, such as autonomous vehicles, facial-recognition check-ins or algorithmic recruiting) and 11 ethical principles (such as “explainability”, “privacy and data rights” or “fairness and non-discrimination”). These lists were developed based on numerous interviews with industry and policy experts, as well as a review of major ethical principles and guidelines produced by leading governmental bodies, non-governmental organizations and big tech companies (for details, see the Appendix). The respondents were also asked for their perspectives regarding the future of AI governance, policy and regulation.

¹ The governance of AI ethics could occur through several mechanisms, including government regulation, technical standards and self-regulation. This report uses the terms “governance”, “policies” and/or “policymaking” to encompass all of these options.

² The policymaker arm of the survey included legislators and regulators, as well as other policy experts who influence policymaking. We use the term “policymaker” in this report for brevity.

KEY SURVEY FINDINGS



Policymakers have a clear vision of AI ethical risks – and are moving to implementation

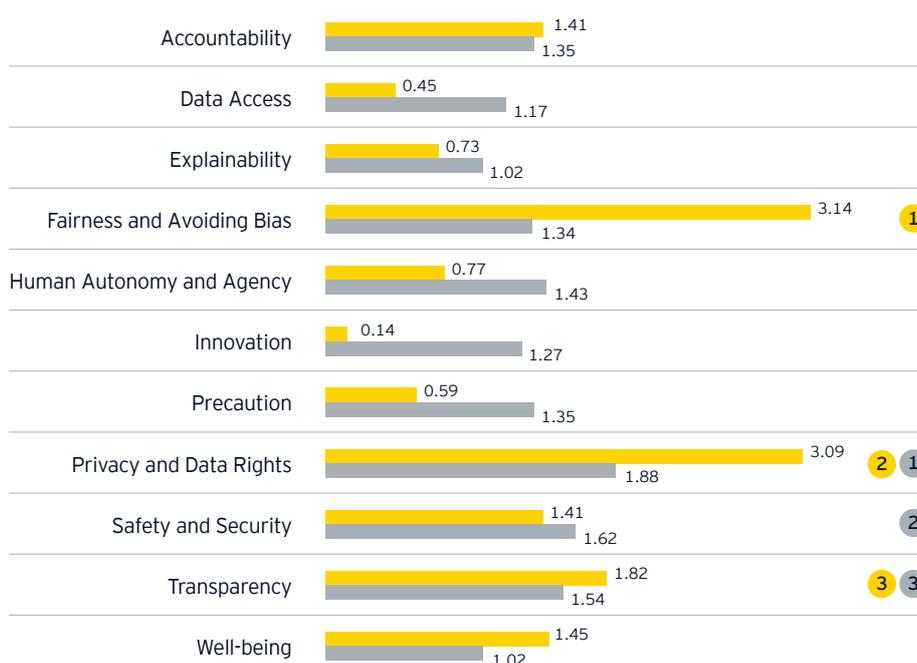
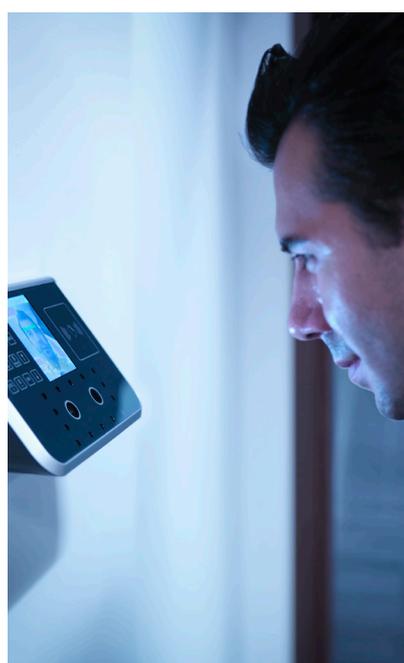
We are at a critical transition point in the governance of AI ethics, as the locus of activity shifts from articulating ethical principles to setting rules and implementing them. Policymakers have achieved consensus on the ethical principles they intend to prioritize and there are signs of increasing collaboration across jurisdictions (e.g., between the European Commission and US).

Policymakers’ strong consensus is evident in the survey data. Policymakers’ responses show widespread agreement on the ethical principles most relevant for different applications of AI. Their responses gravitate toward the one or two specific principles most relevant for each use case and they show strong preferences for these principles.

Consider the use of AI for facial recognition check-ins, in which facial recognition cameras are used to facilitate faster check-ins at establishments such as airports, hotels and banks. Policymakers show a clear ethical vision, rating “fairness and avoiding bias” and “privacy and data rights” as the two most important principles, and doing so by a wide margin. This recognizes the considerable privacy concerns raised by facial recognition technology and the potential for bias (e.g., from misidentification of individuals). On the other hand, companies’ responses were fairly evenly distributed across all ethical principles in an apparently undifferentiated way. (See Exhibit 1).

Exhibit 1. Facial recognition check-ins: policymakers have a clear vision of AI ethical principles

■ Policymakers ■ Companies

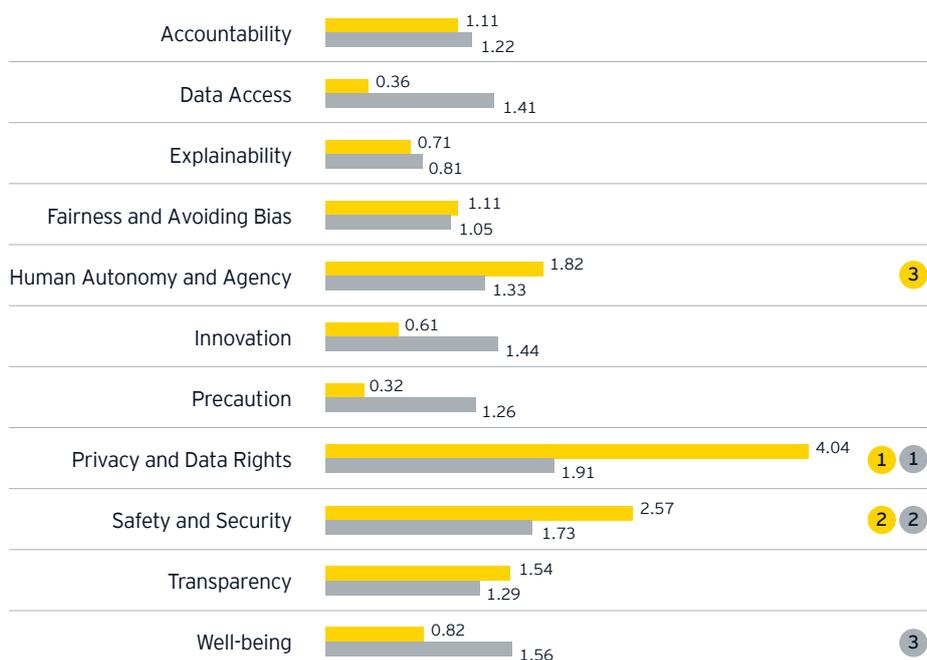


Source: EY and The Future Society (TFS) survey of policymakers and companies. Exhibit shows the importance policymakers and companies assign to various ethical principles in guiding the use of AI for facial recognition check-ins. A higher score indicates greater importance assigned to that ethical principle. For more, refer to the Appendix.

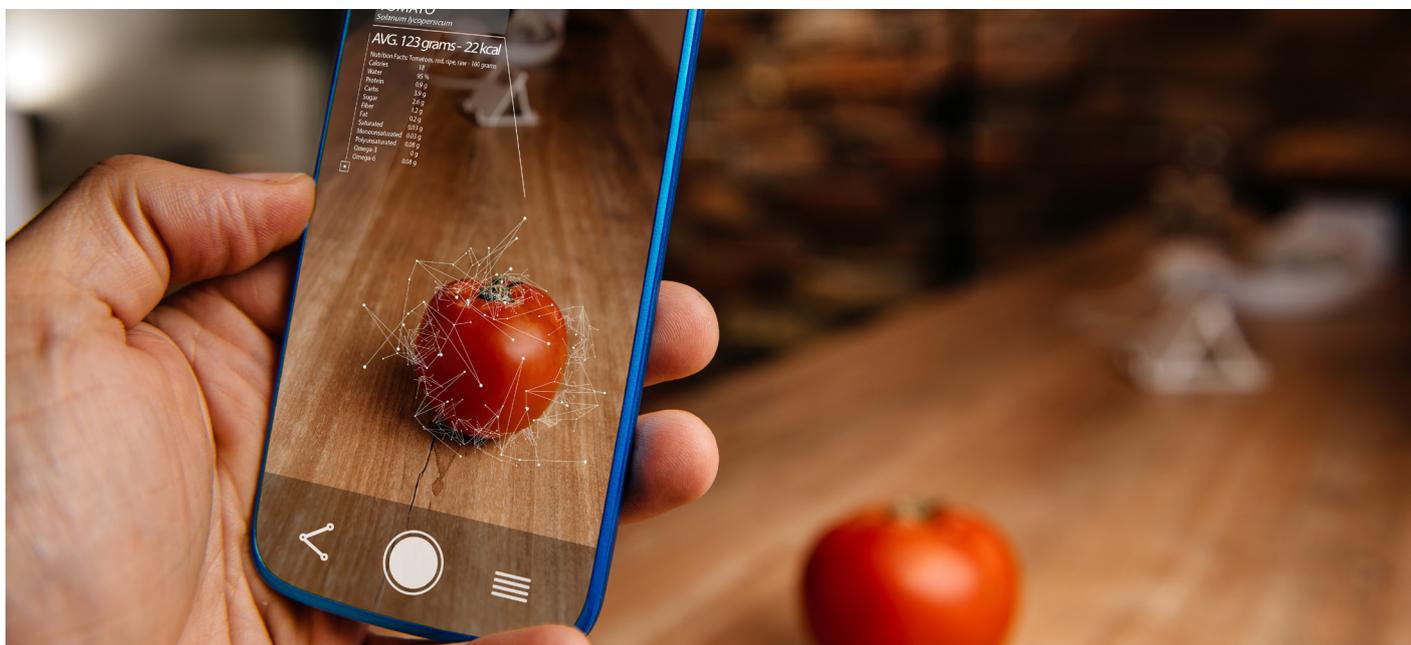
A similar pattern is visible in responses to questions about home virtual voice assistants. Here, policymakers picked “privacy and data rights” as the top concern by a wide margin. Once again, companies’ responses were fairly undifferentiated across the ethical principles. (See Exhibit 2).

Exhibit 2. Home virtual voice assistants: policymakers have a clear vision of AI ethical principles

■ Policymakers ■ Companies



Source: EY and TFS survey of policymakers and companies. Exhibit shows the importance policymakers and companies assign to various ethical principles in guiding the use of AI for home virtual voice assistants. A higher score indicates greater importance assigned to that ethical principle. For more, refer to the Appendix.



Companies are misaligned with policymakers at this critical moment

In contrast, a much weaker consensus exists among companies. Their prioritization of ethical principles is relatively undifferentiated, with evenly distributed responses across use cases and principles. Their top choices are also preferred by narrower margins than those of policymakers.

More importantly, companies are focused on the principles prioritized by existing regulations such as GDPR (e.g., privacy and cybersecurity) rather than on emerging issues that will become critical in the age of AI (e.g., explainability, fairness and non-discrimination).

For instance, companies place “privacy and data rights” among the top three ethical principles in 10 out of 12 use cases and rank it the top concern in seven of those use cases. Policymakers also rate this ethical principle highly, but to a lesser degree – ranking it the most important use case just twice, and instead placing it second or third in nine use cases (see Exhibit 3).

Exhibit 3. Companies rate “privacy and data rights” the most important issue

	Rank assigned by policymakers:			Rank assigned by companies:		
	Most important	Second-most important	Third-most important	Most important	Second-most important	Third-most important
Algorithmic financial planners		✓		✓		
Algorithmic health care providers		✓				✓
Algorithmic recruiting			✓	✓		
Behavioral modification		✓		✓		
Facial recognition check-ins		✓		✓		
Fully autonomous vehicles						
Home virtual voice assistants	✓			✓		
Human emotion analysis	✓				✓	
Information curation and distribution		✓		✓		
Law enforcement surveillance		✓		✓		
Personalized algorithmic pricing		✓				
Social credit and underwriting		✓				✓

Source: EY and TFS survey of policymakers and companies. Chart shows the degree to which companies and policymakers rated “privacy and data rights” among the three most important ethical principles for each use case.

Companies also rate “safety and security” highly, placing it among the top three ethical principles in 10 out of 12 use cases – more than twice as often as policymakers (see Exhibit 4).

Exhibit 4. Companies rate “safety and security” more highly than do policymakers

	Rank assigned by policymakers:			Rank assigned by companies:		
	Most important	Second-most important	Third-most important	Most important	Second-most important	Third-most important
Algorithmic financial planners			✓			
Algorithmic health care providers	✓			✓		
Algorithmic recruiting						✓
Behavioral modification					✓	
Facial recognition check-ins					✓	
Fully autonomous vehicles	✓			✓		
Home virtual voice assistants		✓			✓	
Human emotion analysis				✓		
Information curation and distribution						
Law enforcement surveillance					✓	
Personalized algorithmic pricing					✓	
Social credit and underwriting				✓		

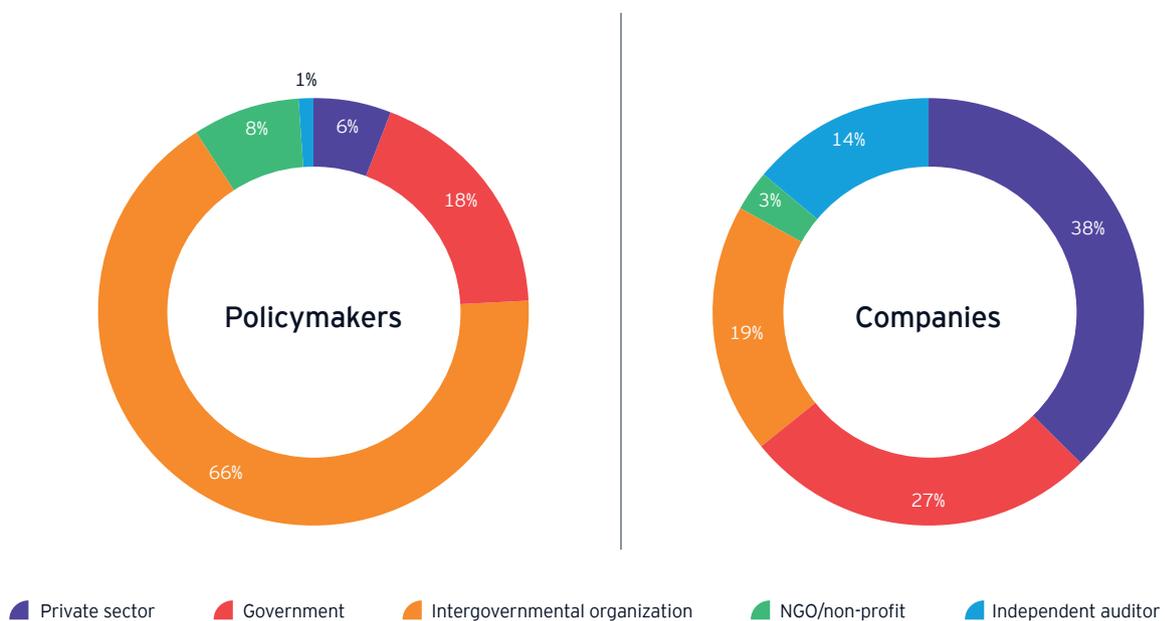
Source: EY and TFS survey of policymakers and companies. Exhibit shows the degree to which companies and policymakers rated “safety and security” among the three most important ethical principles for each use case.

Some additional context about the survey instrument is useful here. “Safety and security” was always listed in survey questions along with a parenthetical explanation: “cybersecurity and safeguarding individuals’ life/health” (see Exhibit A2 in the Appendix). The presence of the word “cybersecurity” in this label suggests companies may have gravitated towards issues they are familiar with – privacy and cybersecurity – and which current laws regulate.

Companies’ focus on currently regulated issues rather than the ethical issues raised by AI may reflect incentives. Companies have a smaller, more narrow set of stakeholders than do policymakers, and their goal is to maximize revenue and financial value. On the other hand, policymakers have a longer time horizon and a more diverse set of stakeholders. Consequently, policymakers tend to focus relatively more on principles that are socially beneficial and less tangible, such as fairness, human autonomy and explainability.

The misalignment between companies and policymakers is also evident in their expectation about the future direction of governance. Given the complex and technical issues at play, many expect that a multi-stakeholder approach will be needed, and that industry has an important role to play. However, policymakers and companies disagree on what form this will take. While 38% of companies expect the private sector to lead a multi-stakeholder framework, only 6% of policymakers agree (instead, two-thirds of them think an intergovernmental organization is most likely to lead).

Exhibit 5. Policymakers and companies disagree on who will lead a multi-stakeholder framework



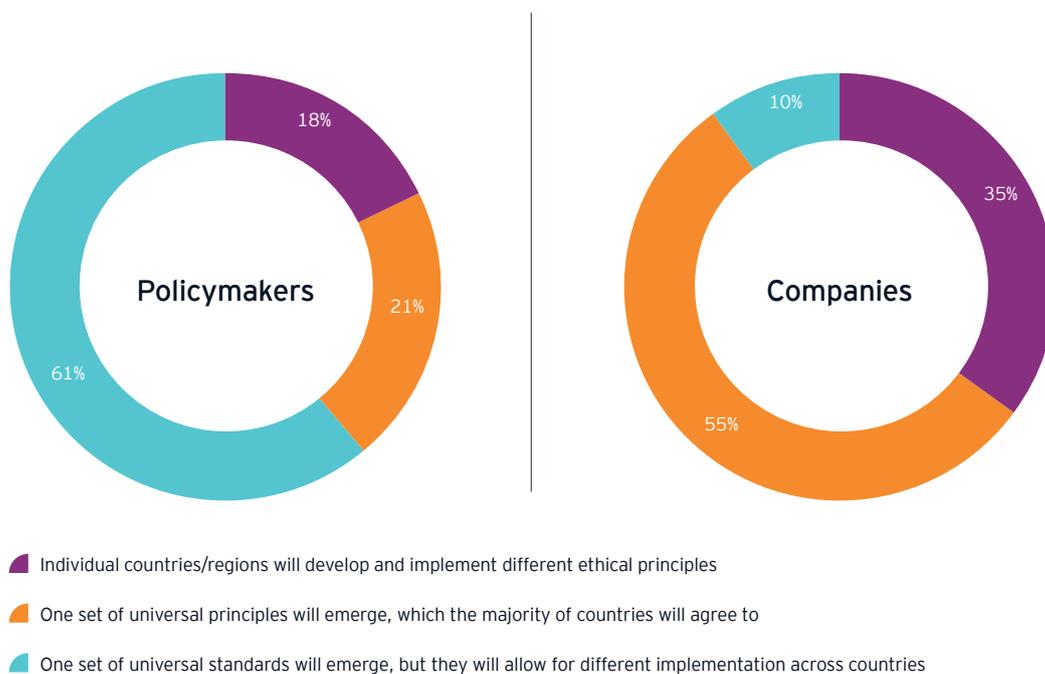
Source: EY and TFS survey of policymakers and companies. Exhibit shows responses to the following question: "Many agree that AI governance will be conducted through a multi-stakeholder framework. Who is most likely to lead this framework?"



Companies' misalignment with policymakers creates new risks

New regulations almost inevitably bring new challenges and risks for companies. Two years after it went into effect, GDPR is still raising a plethora of compliance challenges for companies operating in Europe. Any move to regulate the ethics of AI will likely create new risks for firms. Consider the issue of compliance across multiple jurisdictions. When asked "how much variation will likely occur in developing and implementing AI ethical principles across countries/regions", 55% of companies said they expected "one set of universal principles will emerge, which the majority of countries will agree to" – a fairly uniform outcome which would hopefully make compliance relatively simple. On the other hand, policymakers – who are a lot closer to policy decisions – expect much more complexity and ambiguity, saying that "one set of universal standards will emerge, but they will allow for different implementation across countries." (See Exhibit 6).

Exhibit 6. How much variation will likely occur in developing and implementing AI ethical principles across countries/regions?



Source: EY and TFS survey of policymakers and companies. Exhibit shows responses to the following question: "How much variation will likely occur in developing and implementing AI ethical principles across countries/regions?"

This is part of a larger pattern. Companies' misalignment with policymakers will likely create a host of risks in this space, since firms may be developing products and services that are misaligned with the market and regulatory environment in which they emerge:

- 1. Market or competitive risk.** Consumers are expressing strong concerns specific to different applications of AI, and policymakers are proceeding to address these concerns. Firms whose products or services don't address these concerns will have fundamentally misread market demand, and risk losing market share.
- 2. Reputational risk.** Today's consumers have big megaphones thanks to social media. AI products and services that cause damage to a consumer could erode the company's brand and perception.
- 3. Compliance risk.** If companies are not actively involved in shaping emerging regulations and don't understand the ethical principles policymakers are prioritizing, they risk developing products and services that aren't designed to comply with future regulatory requirements.
- 4. Legal risk.** The inability to comply with regulations, could open companies to litigation and financial penalties.

To mitigate these risks, it is in companies' best interest to coordinate with policymakers in developing realistic and effective policy measures and governance frameworks.

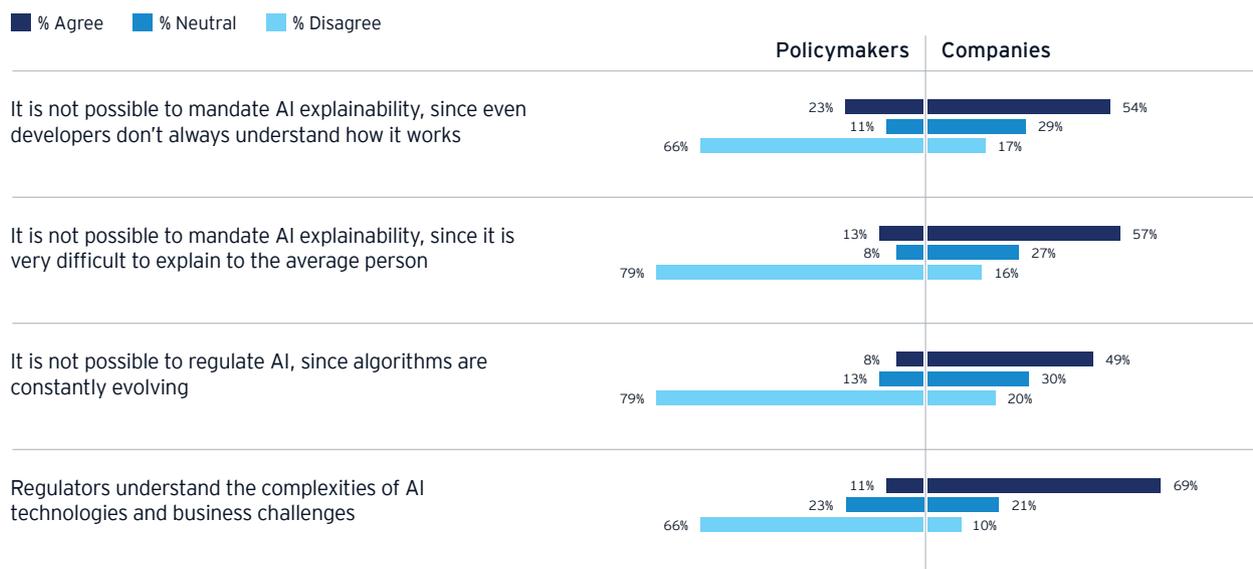


Companies and policymakers have contrasting gaps

Policymakers and companies have gaps that are mirror images of each other.

While policymakers understand the “big picture” ethical concerns raised by AI applications, the survey suggests that they might not be as immersed in technical and business details as companies. Two-thirds of policymakers surveyed indicate “regulators don’t understand the complexities of AI technologies and business challenges.” Policymakers’ responses also indicate they are less sympathetic to technical or business challenges in the implementation of ethical principles. For instance, 79% of policymakers disagree with the statement that “it is not possible to mandate AI explainability since it is very difficult to explain to the average person”. An identical percentage of policymakers disagree that “it is not possible to regulate AI, since algorithms are constantly evolving” and two-thirds of them disagree that “it is not possible to mandate AI explainability since even developers don’t always understand how it works.” (See Exhibit 7).

Exhibit 7. Policymakers are less immersed in technical and business challenges



Source: EY and TFS survey of policymakers and companies. “% Agree” shows the share of respondents who either agreed or strongly agreed with each statement. “% Disagree” shows the share of respondents who either disagreed or strongly disagreed with each statement.

While companies are less aligned with the salient ethical principles governing AI, they presumably have greater depth on business and technical issues.

Admittedly, assessing the technical skills of policymakers was not a major focus of the survey, and these inferences are based on three questions. Nonetheless, the responses are directionally consistent, are based on self-reported data, and suggest that both stakeholders have potential blind spots – recognition of ethical issues for companies and deep technical or business expertise for policymakers.

These contrasting gaps require a comprehensive approach to regulation

How do we bridge these gaps between policymakers and companies? The most obvious answer is through closer collaboration – for instance, policymakers using a consultative and deliberative approach, with input from the private sector, especially on technical and business complexities.

More fundamentally, the survey results suggest the need for a comprehensive approach to regulation with policymaker-led and company-led components, since both parties have strengths that only partially cover the domain knowledge required. This would include:

- ▶ **Hard regulation:** traditional regulatory approaches created and implemented by governmental bodies
- ▶ **Soft regulation:** standards and best practices developed by trade associations, professional organizations or communities of interest
- ▶ **Self-regulation:** principles and codes of conduct developed and implemented by individual companies

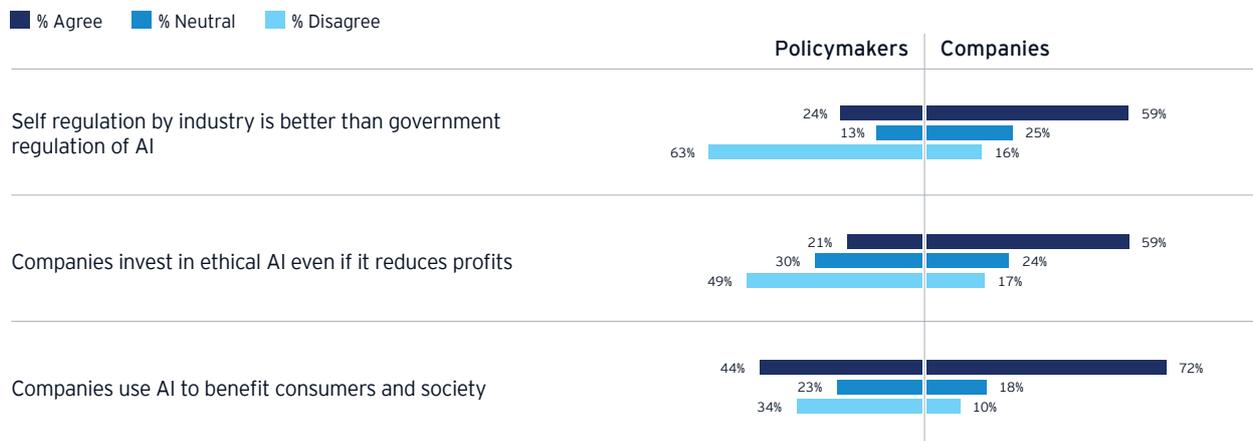


A trust gap complicates matters – and creates an opening for third parties

However, collaboration will itself be challenging because of another gap – a trust deficit between companies and policymakers.

Policymakers don't trust the intentions of companies. Almost six in ten company respondents (59%) agree that "self-regulation by industry is better than government regulation of AI" while a similar share (63%) of policymakers disagree. Furthermore, while 59% of company respondents agree that "companies invest in ethical AI even if it reduces profits", only 21% of policymakers agree (49% disagree). Meanwhile, 72% of companies agree that "companies use AI to benefit consumers and society", while only 44% of policymakers share that view. (See Exhibit 8).

Exhibit 8. Companies have a trust gap with policymakers

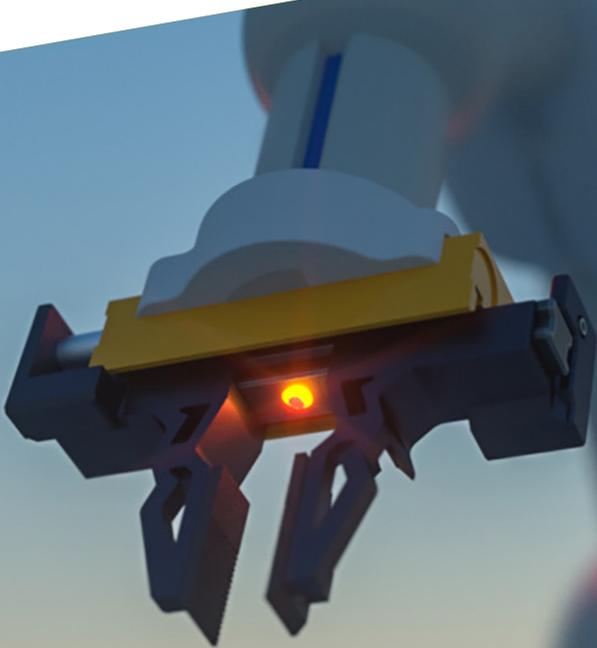


Source: EY and TFS survey of policymakers and companies. "% Agree" shows the share of respondents who either agreed or strongly agreed with each statement. "% Disagree" shows the share of respondents who either disagreed or strongly disagreed with each statement.

Bridging this trust gap will not be easy. Trust is easy to lose; it is more difficult and time-consuming to regain. For tech companies, trust has been eroding for a while now because of the techlash. One step that would certainly help is for firms to recognize the fears that new applications of AI are generating, which would demonstrate to consumers and policymakers that they are not out of touch with the most important ethical issues.

Independent third parties could play a constructive role here. Consider the “mirror image” nature of the gaps between policymakers and companies, with contrasting strengths and weaknesses in ethical issues and technical/business expertise. Third parties could be perfectly positioned to bridge these gaps. An objective third party could be more trusted than tech companies while better understanding technical complexities than policymakers. And, to an uneasy public, a third party could offer assurance by ensuring compliance with regulations and providing transparency and simplification – creating the equivalent of a “Fairtrade ” or an “Organic” label for AI.

A comprehensive approach to regulation, coupled with a critical bridging role played by independent third parties, could create an “infrastructure of trust”. This infrastructure of trust would allow policymakers to protect public safety, while allowing businesses to develop and deploy AI in ways that benefit consumers.



Implications for companies

1. Focus on AI's emerging ethical issues. GDPR was just the beginning. AI will raise a host of new ethical challenges.

- ▶ Which AI ethical principles are most important in your sector or segment?
- ▶ How will they affect your business?
- ▶ How should your strategy respond?

2. Engage with policymakers. If you're not at the table, you're on the menu. Policymakers are ready to move ahead – but, without industry input, blind spots could lead to unrealistic or onerous regulations.

- ▶ How do policymakers view AI governance and regulation in your sector or segment?
- ▶ What real-world issues are critical to understanding your business?
- ▶ How will you be part of the conversation?

3. Be proactive with soft and self-regulation. Stakeholders expect more now. If companies want to lead on AI innovation, they need to lead on AI ethics as well.

- ▶ Have you developed a corporate code of conduct for AI – and does it have teeth?
- ▶ How aligned is it with the ethical principles consumers and policymakers prioritize?
- ▶ How are you working with your peers (e.g., through trade organizations) on these issues?

4. Understand and mitigate risks. AI governance – and particularly the “hard regulation” variant – will create new challenges and risks. Companies' misalignment with policymakers only increases those risks.

- ▶ What risks might the move to AI governance/regulation create for your business?
- ▶ How are you mitigating and preparing for these challenges?

Implications for policymakers

1. Consult and deliberate. We need approaches that align stakeholders' interests and are technically practical. Open consultation with a range of stakeholders is key.

- ▶ How are you engaging with technical and industry experts to understand implications across the range of AI use cases?
- ▶ How are you hearing the critical voice of the citizen on sensitive AI applications?

2. Proceed with appropriate speed. AI is moving fast and ethical concerns are real. Policymakers need to move quickly – but also carefully, with flexible and adaptive approaches.

- ▶ What processes do you have for monitoring fast-changing technologies and markets?
- ▶ How are you developing policies and regulations that are adaptive to changing conditions?

3. Align globally. International coordination is needed to tackle issues consistently, mitigate global risks, and learn from leading countries.

- ▶ How are you coordinating with policy and industry partners across borders to manage risks and harmonize approaches?
- ▶ What could you learn from other countries?

Appendix

METHODOLOGY

Survey methodology

The survey covered 12 “use cases” (i.e., applications of AI, listed in Exhibit A1) and 11 ethical principles (Exhibit A2). Use cases, ethical principles and statements were developed based on interviews with leading AI industry, policy and academic experts. The process was also informed by a literature review of ethical guidelines published by several organizations, including: AI4People, Asilomar AI Principles, Beijing AI Principles, European Commission (High-Level Expert Group on AI), Google, IEEE (Ethically Aligned Design v2), Microsoft, Montreal Declaration for Responsible AI, and OECD AI Principles.

The survey was web-based and conducted in late 2019 and early 2020 by Thought Leadership Consulting, the strategic content division of Euromoney Institutional Investor Plc. It obtained responses from 71 policymakers and 284 companies across 55 countries. The larger number of company responses reflects that the survey covers use cases that are mostly industry-specific (e.g., autonomous vehicles, algorithmic financial planners, home virtual assistants). So, the typical company respondent was able to answer questions about fewer use cases than the average policymaker and more responses were required to achieve a sufficient response rate in each individual use case. Moreover, AI policy being a nascent field, a smaller pool of AI policymakers and policy experts were available to complete the survey. Two use cases (“product delivery drones” and “employee productivity tracking”) were dropped because of insufficient response rates from both policymakers and companies, bringing the initial list of 14 down to 12.

Only fully completed survey responses were included in the analysis. An average of 25 responses per use case was achieved for the policymakers' survey, from a minimum of 20 responses for “Information curation and distribution” and “fully autonomous vehicles” to a high of 33 responses for “algorithmic recruiting”. For the companies' survey, the average sample achieved was 92 responses per use case.

To ensure the quality of responses, both sets of respondents were self-screened for knowledge about AI and its ethical risks. Company respondents who reported low levels of knowledge and whose companies were not actively implementing AI for a particular use case were not shown questions pertaining to that use case. Similarly, policymakers were asked to select the use cases in which they were most knowledgeable and were only shown questions pertaining to those use cases.

The use cases and ethical principles were displayed to individual respondents in a randomized order to reduce the potential for response bias.

For each relevant use case, respondents were asked to pick the five ethical principles they consider most important, (ranked from “most important” to “fifth-most important”). For some of the summary exhibits, these selections were assigned a score of 5 (most important) to 1 (fifth-most important), and individual respondents' scores for each ethical principle were averaged to obtain an aggregate score. This average score was weighted by the total number of respondents for the relevant use case.

The survey also asked questions to measure company and policymaker sentiment around the use, risks and regulation of AI. This included questions about the degree to which respondents agreed or disagreed with a series of statements on the role of appropriate role of policy, regulation and business. It also included questions about the likely future direction of AI governance.

Exhibit A1. Use cases

	Use case	Description
1	Algorithmic health care providers	AI makes diagnoses and prescribes medical interventions
2	Behavioral modification	AI “nudges” people in highly customized ways to change behaviors (e.g., health, saving/spending)
3	Home virtual voice assistants	Voice assistants that help residents (elders, children, adults etc.) in the home on a variety of tasks
4	Social credit and underwriting	AI uses big data to get 360-degree view of individuals’ behaviors and make decisions about loan applications, premium pricing, etc.
5	Algorithmic financial planners (Robo-advisory)	AI agents monitor personal finances and conduct financial/retirement planning on behalf of clients
6	Fully autonomous vehicles	Vehicles that drive themselves without human input (Level 4 or 5)
7	Facial recognition check-ins	Facial recognition speeds check in at airports, hotels, banks, etc.
8	Law enforcement surveillance	Widespread use of facial recognition, license plate scanners, drones etc. by law enforcement and security services
9	Personalized algorithmic pricing	Algorithms change prices for individual consumers based on data about their habits, preferences and circumstances
10	Algorithmic recruiting	Using algorithms to screen applicants and make hiring decisions
11	Information curation and distribution	Algorithms curate, analyze, prioritize and amplify digital content (algorithms that pick and choose what news/information to show you)
12	Human emotion analysis	Sensors and computer vision combined with algorithms that predict and track human emotions

Exhibit A2. Ethical principles

	Ethical principle	Definition
Well-being, sustainable development and inclusive growth		
1	Well-being (sustainability, human rights, equitable benefit)	This principle emphasizes that AI use should not infringe on internationally recognized human rights. Developers should design to maximize societal well-being, enhance sustainability, and ensure that access and benefits are distributed equitably.
2	Innovation (encouraging innovation)	This principle prioritizes deploying AI innovation quickly. Regulators should not hinder innovation, based on the assumption that innovation benefits individuals.
3	Data access (open access to AI datasets)	This principle emphasizes the need for providing open access to data, through mechanisms such as public datasets, open data platforms, data marketplaces and data trusts.
Prevention of harm		
4	Precaution (moving carefully in poorly understood areas)	This principle emphasizes the need to anticipate and address catastrophic or existential risks, since advanced AI could profoundly change human existence.
5	Safety and security (cybersecurity and safeguarding individuals' life/health)	This principle states that AI applications should not endanger the lives of individuals or produce adverse impacts on their health. They should include safeguards against hacking and other criminal misuse.
Human-centered values and fairness		
6	Privacy and data rights (individuals' privacy and control over their data)	This principle asserts that people should have the right to access and control the data they generate. AI should protect the privacy of individuals, including through data anonymization.
7	Human autonomy and agency (retaining human control)	This principle prioritizes self-determination and free will. Humans should choose how and whether to delegate decisions to AI, to accomplish human-chosen objectives. AI should promote users' preferences, and set limits for system intervention.
8	Fairness and avoiding bias (nondiscrimination)	This principle states that AI applications should minimize the risk of results that are (intentionally or unintentionally) biased or discriminatory.
Explicability and accountability		
9	Explainability (ability to explain decision-making process)	This principle states that individuals have the right to understand how an AI decision was made. Explanations should be comprehensible by the "person on the street."
10	Transparency (divulging AI usage and limitations)	This principle asserts that the use of AI should be disclosed, as should the capabilities and limitations of AI applications. (Note that, unlike the Explainability principle, this is not about explaining how decisions are made, but is instead about disclosing whether AI is being used at all.)
11	Accountability (auditability, fault-finding mechanisms, penalties)	This principle requires that relevant parties should be accountable for any negative consequences. This includes mechanisms such as fault finding, monetary compensation and the right to appeal (contestability).

Contacts

Executive sponsors

Gil Forer

EY Global Markets Digital and
Business Disruption Leader, EYQ leader
gil.forer@ey.com

Nicolas Mialhe

The Future Society, Founder & President
nicolas.mialhe@thefuturesociety.org

Research and content team | EY

Ben Falk

EYQ Insights Director
ben.falk@uk.ey.com

Gautam Jaggi

EYQ Insights Director
gautam.jaggi@ey.com

Prianka Srinivasan

EYQ Insights Director
prianka.srinivasan@ey.com

Research and content team | The Future Society

Sacha Alanoca

Senior AI Policy Researcher
sacha.alanoca@thefuturesociety.org

Adriana Bora

AI Policy Researcher
adriana.bora@thefuturesociety.org

Arohi Jain

Senior AI Policy Advisor
arohi.jain@thefuturesociety.org

Yolanda Lannquist

Head of Research & Advisory
yolanda.lannquist@thefuturesociety.org

About EY

EY is a global leader in assurance, tax, transaction and advisory services. The insights and quality services we deliver help build trust and confidence in the capital markets and in economies the world over. We develop outstanding leaders who team to deliver on our promises to all of our stakeholders. In so doing, we play a critical role in building a better working world for our people, for our clients and for our communities.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. For more information about our organization, please visit ey.com.

About EYQ

EYQ, EY's global think tank, generates new insights by bringing together business, the public sector and academia to challenge entrenched thinking, shift perceptions and help catalyze change. We want to sense new trends early, and understand and communicate their implications quickly and powerfully. By seeking the answer to "What's after what's next?", we help leaders anticipate the forces shaping our future – empowering them to seize the upside of disruption and build a better working world.

© 2020 EYGM Limited.
All Rights Reserved.

EYG no. 004607-20Gb1
1912-3340731
ED None

This material has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax or other professional advice. Please refer to your advisors for specific advice.

ey.com

About The Future Society

The Future Society is an independent 501(c)(3) nonprofit think-and-do tank. Specialized in questions of impact and governance, their mission is to help advance the responsible adoption of AI and other emerging technologies for the benefit of humanity. The Future Society leverages a global, multidisciplinary network of experts, practitioners, and institutional partners. Their work encompasses developing policy research and advisory services, seminars and summits, educational and leadership development programs, as well as special projects. The Future Society is independent yet connected to all relevant stakeholders, and tackles a broad, but carefully selected, range of short-term and longer-term issues in AI governance. The Future Society's expertise, independence and commitment to advancing the governance and trustworthiness of AI have earned them the trust of key stakeholders in the ecosystem including governments, industry, international organizations, academia and civil society - across geographies. The Future Society is listed among the "Best Artificial Intelligence Think Tanks 2019" by the University of Pennsylvania Lauder Institute's Global Go To Think Tank Index.