Atlantic Council
GEOTECH CENTER

Policy on a Spectrum: Guiding Technology Regulation Through Value Tradeoffs

Steven Tiell
Lara Pesce Ares

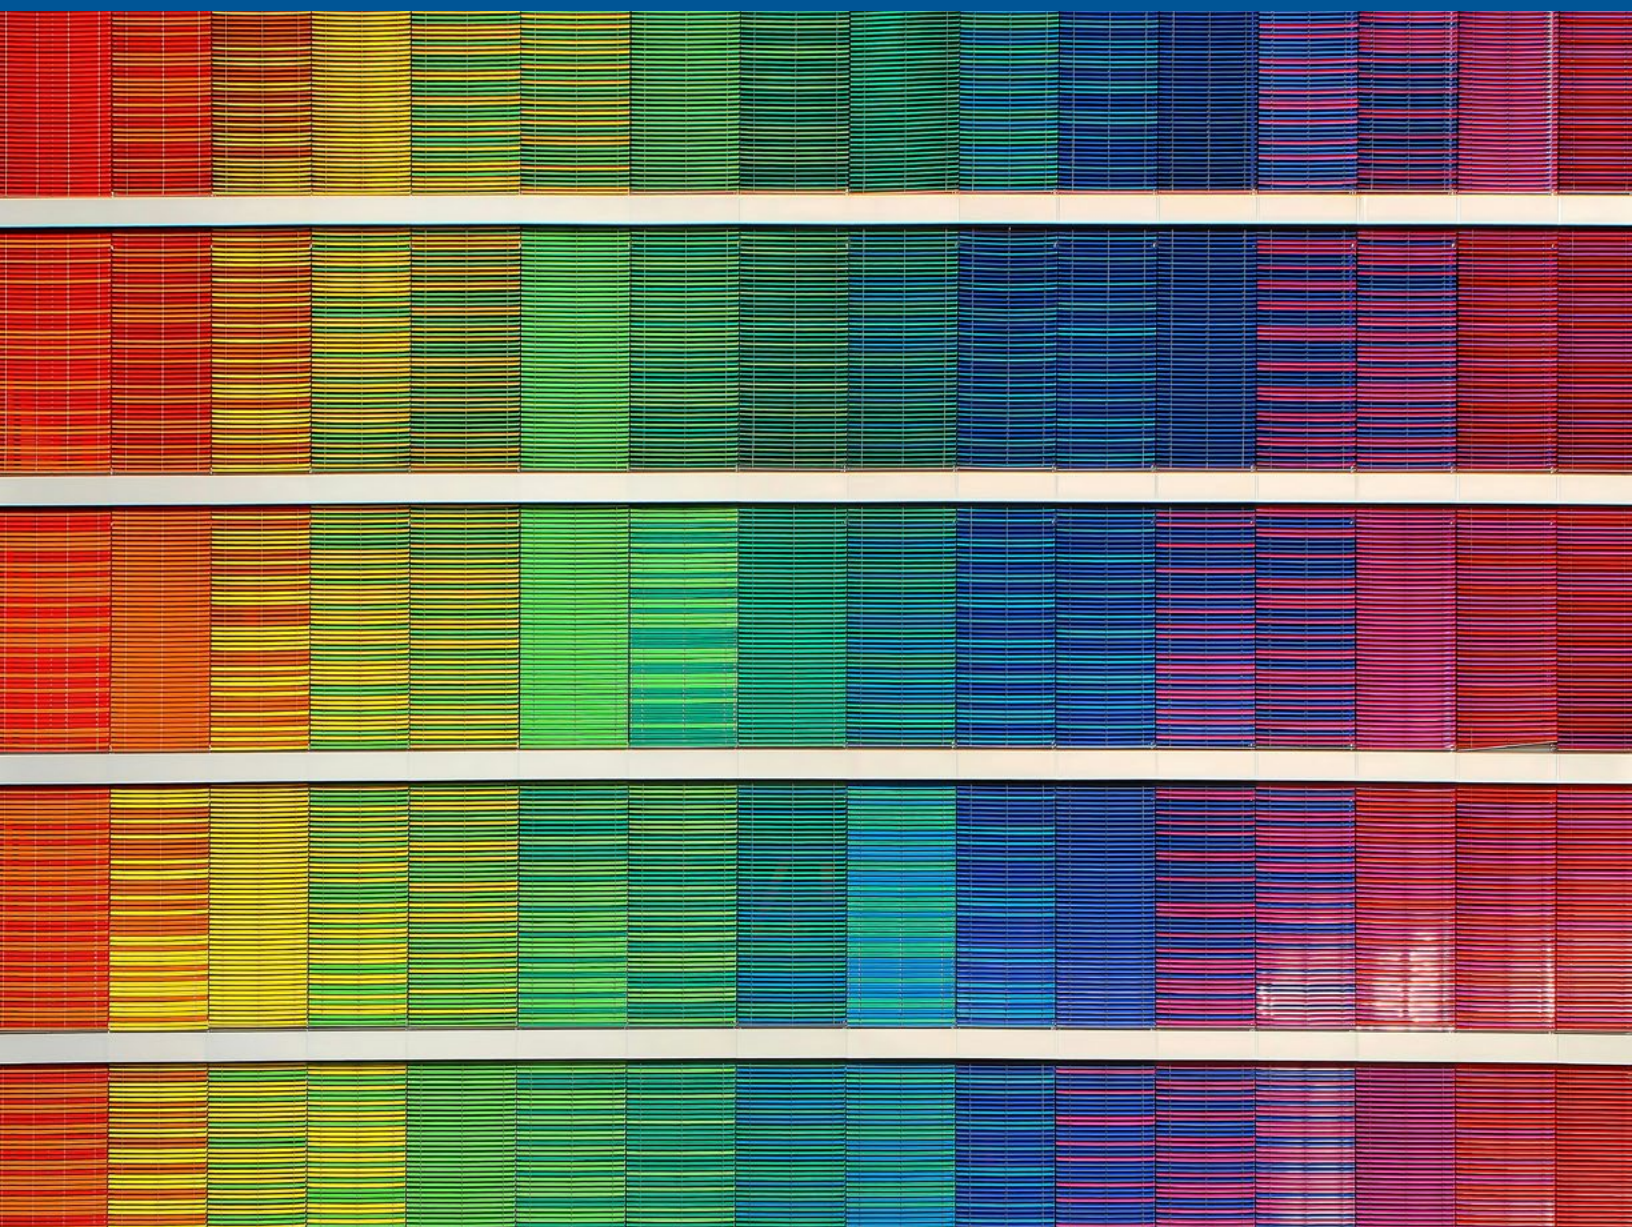

COVER AND INTERIOR IMAGES: Unsplash

ISBN-13: 978-1-61977-311-0

March 2024

Policy on a Spectrum: Guiding Technology Regulation Through Value Tradeoffs

Steven Tiell
Lara Pesce Ares

TABLE OF CONTENTS

I. INTRODUCTION

We live in an era marked by rapid digital transformations, a time when every facet of our culture—including businesses, governments, and civil societies—is undergoing an unyielding technological metamorphosis. The relentless, ever-increasing pace of this change presents a stark paradox: while societal norms and behaviors shift to accommodate this new digital reality, our legal and regulatory systems often find themselves in a perpetual game of catch-up from a position further and further behind. More recently, major regulatory frameworks have employed an arguably more effective risk-based approach, in which requirements can become more stringent based on risk categorization. However, these regulations are limited in jurisdiction, and will need to be updated over time as new technologies emerge. Furthermore, they aim to set the bounds for allowable behavior, and do not prescribe specific actions for organizations in most contexts. In this underregulated environment, private organizations, relying only on basic compliance to restrict decision-making, are introducing systemic risks that could jeopardize the fabric of societies. Action must be taken by organizations in the public and private spheres for the continued viability of the commons and for the benefit of all.

In response to this adaptive dilemma, a dynamic and proactive framework can help to inform and guide technology policy decisions. This framework is designed primarily for technology policymakers, as well as leaders in the private sector who are keen on setting and following best practices that anticipate, and exceed, existing compliance standards. Additionally, developers, designers, and business leaders can leverage this framework to reconcile business imperatives with ethical and societal concerns, thereby mitigating risk and building trust in their technology.

Numerous frameworks, such as those from the European Union (EU), the US National Institute of Standards and Technology (NIST), the Executive Office of the President under the Biden-Harris administration, and the Organisation for Economic Co-operation and Development (OECD), have tried to tackle technology policy challenges.[1] These efforts need to continue at pace if societies are to withstand the barrage of systemic risks from advanced technologies—and the framework discussed herein will contribute to augmenting transparency, ethics, and rigor in modern technology policy. This framework positions policies along a spectrum between two competing values, indicating which value carries more weight in the resulting policy. It is distinctly different from—and an additive to—existing regulatory frameworks.

This approach aims to provide a nuanced perspective, stimulate discussions, and delineate guardrails for teams responsible for creating and upholding requirements for the design, development, and governance of technology. It offers a broader context for regulatory precedent and encourages forward-thinking policy decisions. The private-sector audience will be prompted with a proactive, risk-based approach that encourages entities to go beyond the legally required minimum, or to be "forward compliant." For them, this will entail anticipating, preparing, and implementing controls for future legislation, thereby fostering a culture of proactive compliance.

This idea can be represented on the following spectrum.

Proactive risk-based approach > Legally required minimum

Risk-based decision-making proactively protects both consumers and the organization.

Risk-based decision-making is consistent with a desire to achieve "forward compliance," a proactive approach to current risks and potential forthcoming regulation. For example, in creating internal privacy policies, it is reasonable to assume a regulator will legislate around personally identifiable information (PII), and organizations would strive to be compliant. A risk-based privacy approach goes further and applies protections greater than mere compliance, perhaps extending to include PII inferences and increasing measures of security beyond the minimum requirement. This approach also addresses risks that arise over time, especially as business practices and technologies evolve.[2]

For a deeper exploration of the value spectrums framework itself, please review Principles to Practice: Using Ethical Spectrums to Guide Decision-Making, the companion paper in which these spectrums were originally introduced.[3] This paper employs the same sections and spectrums as its predecessor, but has adjusted the order and articulation to increase public-sector relevance.

Throughout the policy spectrums, policy examples from major technology regulatory regimes will be referenced, showing where they land along the policy spectrum.

These examples highlight where the policy recommendations align with major current regulatory frameworks. The example coverage is representative of the current regulatory landscape, which is largely led by European legislation. Both private- and public-sector readers may find it useful to note where there are gaps in regulatory coverage, as those may be opportunities to advance current standards.

In the private sector, there is an opportunity to use the framework to ensure the decision-making in the organization reflects its value priorities. Wherever an organization chooses to be along any one of these spectrums, the deliberate process of weighing priorities will help to connect product features directly to organizational values, reinforcing culture and optimizing for distributed decision-making. This process also helps to substantiate policy decisions that might be communicated publicly.[4]

In this framework, the overarching goal is to protect the continued sustainability and enrichment of the human condition. This notion is further extended to technology policy decision-makers, who set the outer bounds of possibility for each spectrum with a bias toward protecting a baseline quality of life for individuals.

In summary, this framework aims to balance the profound role of technology in advancing societies with the collective interest in ensuring the safety and security of individuals and groups. By integrating technology advancements with conscientious regulatory foresight, these proposals are intended to walk a fine line of continuing to embrace innovation while stepping up efforts to protect the most vulnerable—with ambition toward a sustainable, inclusive, and secure digital future.

II. DATA PROCUREMENT AND USE

Local cultural norms can be widely variable and for data (and its descendant uses) to be valuable, data must meet the local standards and stakeholder expectations. To achieve this, organizations and institutions might take a risk-, principles-, or rights-based approach to maximize local resonance and minimize harm. The extent an organization might be willing to go will depend upon the potential value from a particular market.

Thoughtful policy design does more than simply protect individuals and organizations through risk mitigation; it has the inherent potential to generate new value by improving relationships and retention with existing stakeholders, while attracting new stakeholders.

The spectrums recommended as a starting point include the following.

Collect relevant data > Collect anything and/or everything possible

Less data may result in both better analysis and less risk.

It is always best to first consider the strategic questions that need to be answered and then figure out how those answers should be informed by data. In some cases, the data may serve multiple purposes, for both the user and the collector. For example, an app could collect location data to give users more relevant search results, but also to personalize ads. It is important to understand all the ways in which the data will be leveraged from the start. After questions and purposes are articulated, data maps can be created to specify the data that must be collected. Then, data scientists can consider data-minimization techniques to further reduce the data needed to answer the questions, while ensuring that the resulting dataset remains adequately representative of the circumstances and populations it is intended to cover. This minimizes the data burden—the infrastructure, processes, and personnel required to handle large volumes of data. This leaves the organization in a strong strategic position, having derived valuable insight with minimal data risk should a breach or leakage occur.

Informed consensual use of data > Exploratory use

Plan for how to use data, be transparent about its use, and gain consent.

The more specific and informed the consent regime, the lesser the future liability and the stronger the trust relationship with the data provider. Data subjects hold a range of expectations about the privacy of data they share and what constitutes acceptable secondary and tertiary uses. These expectations are often context dependent. Designers and data professionals should give due consideration to those expectations, and align products and services accordingly.

Data expiration > Digital perpetuity

Outdated data is a risk to model integrity, informed decision-making, and legal liability.

It might be a priority to keep data as a record or as a resource for future use. However, the longer data is kept, the higher the security and privacy risks become, all while value and public perception are degraded. All data has a useful life. Designers and policymakers should consider this as part of security protocols, consent regimes, and policymaking.

In each section, the recommended spectrums will be applied to specific policies to demonstrate how a policy can vary depending on where it falls along the value tradeoffs. For example, in this "Data Procurement and Use" section, we will cover how opt-in policies, consent policies, and provenance policies vary along these spectrums.

Under each policy matter, the spectrums will continue to be used, with the most aggressive ethical stance on the left and the minimum bar on the right. These areas will employ the continued use of ">" representing "over" as each element is compared to the others, as borrowed from the agile manifesto and proven in practice to align with existing processes. Each will then be described in order from the weakest—that is, the legally required minimum (in any jurisdiction)—to the strongest ethical stance.

Each subtopic discusses a set of must-have considerations for policy-oriented societies to embrace artificial intelligence and the innovations made possible because of it.

OPT-IN AND OPT-OUT POLICIES

An opt-in policy for a particular act of data collection, processing, or storage represents the specific action required by the user to consent. In other situations, where consent can be assumed, opt-out policies become paramount. Opt-in and opt-out policies vary centrally on whether consent is assumed as a default and the extent to which that consent is applied beyond its original context.

Data collection opt-out > Progressive opt-ins > Verifiable opt-in > Non-transferrable opt-in

- **Opt-ins are non-transferrable**. If data is disclosed for one or multiple agreed purposes and the data collector wants to use it for another purpose, they must gain consent from the data disclosers. Ideally, agreed purposes provide reasonable flexibility to the collector, while being specific in nature.

> The General Data Protection Regulation (GDPR) [Article 17] requires data controllers to notify users if they intend to process personal data for a purpose other than that for which the personal data was collected and specifically consented to. Users may exercise their right to opt out of data processing at any time.

- Make **opt-ins verifiable**. The method and context of consent, mechanisms used to grant consent, response(s), and date/time should all be included in metadata. Data professionals should strive to use data in ways that are consistent with the intentions and understanding of the disclosing party.
  - To be able to audit for veracity of consent, even more information is necessary to ensure the veracity of the method and the capacity of the consenting party. For example, establishing lawful bases of consent, age verification, avoidance of consent fatigue, testing of consent interface, tracking of chain of custody, etc.
- Make **opt-ins progressive**. Allow users to proceed anonymously or as a guest immediately, or as early as possible. If users want to do something that requires PII to function, only then should they be asked to volunteer additional information that is strictly relevant to the action at hand.
  - Offer users self (on their own device) or third-party options to store that information and give options to expunge shared PII after the transaction.
  - Give users a series of data-collection options as dynamic as the range of functions they need. For example, requests of location data on mobile devices should prompt users to "allow location data for this instance," instead of presenting them with an option to leave it off or turn it on in perpetuity. This reduces the permanence of data and further aligns data collection with the values and privacy expectations of users.
- Allow users to **opt-out of (or refuse to opt-in to) collection** of PII and/or interaction data while still being able to use the platform—with or without a reasonable fee. Organizations should strive to allow everyone to have access to the social and economic benefits of data, especially data about themselves.

> The California Consumer Privacy Act (CCPA) [1798.135] provides users with the right to opt out of data processing about them. GDPR [Article 6] requires users to opt in before any processing of data about them can occur.

  - Opt-outs should be retroactive. Allow users to opt out at any point and retract all the PII collected up to that point. This may trigger a refactoring of a trained model if the PII removed was used to train the original model.

GDPR [Article 17] stipulates individuals have the right to have personal data erased. This is also known as the "right to be forgotten."

- At any point, users should be able to download all the data a platform might have amassed about them to gain transparency into the collection and use that has occurred up to that point.

GDPR [Article 15] gives data subjects a "right of access," which says that, when requested, any company must provide the user with the personal data about them. The company should provide it in a way that is easy to read, do so in a timely manner, and include background information on how it got the data and how it uses it.

- At any point, users should be able to review and un/verify the data a platform holds about them, particularly when the data is used to make decisions about their experience on the platform or in real life.

CONSENT

Maximizing transparency at the point of data collection can minimize more significant risks later. Design practices that incorporate deliberate decisions about transparency, configurability, accountability, and auditability will serve to mitigate downstream digital risks. To this aim, consent practices should strive to inform the user of what they are consenting to, and what possible consequences they may face.

Delegate consent > Educated consent > Clarity of consent

- Present **a clear list of items to which a user is consenting** regarding the data they are disclosing.
  - To help recipients make informed decisions, organizations should present materials in a way that is approachable to average individuals by focusing on where in the process the consent requests are presented and how they are used. This might mean having a readable, user-friendly version presented prior to the legally binding language. It could also mean listing the consent requests, ranked from most concerning for informed users to least concerning. In industries dealing with highly sensitive data (e.g., health, financial), it is especially relevant to include how data will be protected and the higher risks that are mitigated.
  - Have transparency regarding what questions the data being collected will be used to answer. Digital forms often offer this feature with a "more information" icon that can be hovered over for more details.
- There are circumstances in which interactions with AI could alter "life trajectories" beyond a single instance, such as when employers screen candidates using third-party vendors that then also permanently disqualify them from a number of other employers based on that single screen.[5] These circumstances should require **"educated consent,"** not merely "informed consent," for users to truly understand the technology to which they are consenting and its implications for their lives.

GDPR [Article 7] requires users to give demonstrable consent for personal data processing to occur and obligates the data controller to provide key information including "the existence of any automated decision-making, including profiling...and meaningful information about the logic involved, as well as the significance and the envisaged consequences of such processing for the data subject."

- Give the ability to **delegate consent** requests to a proxy (a human, algorithmic, or hybrid entity that makes privacy- or judgment-oriented choices that someone else might subscribe) after an initial opt-in, removing the burden on the user for making repeated and potentially less informed technical judgments, and allowing this to be a potential service offering for consumers. This recommendation is conditional on the delegation decision being transparent to the user and the user having the ability to retract any delegation decision—or the proxy itself—at any time.
  - On a platform or device, allow users to elect/follow a proxy for their consent, privacy, and/or security settings. This enables users to delegate decision authority to a more informed party (with permission). Examples of proxies include a "super user" with followers for their settings, or predefined setting profiles (such as "less personalized" to "most personalized") that users select based on their preference.

Retain access > Consent to transfer > Opportunity to withhold > Notice of transfer

- Each time data is sold or transferred to another party with separate governance, **notice of that transfer** must be given to the original data discloser.

GDPR [Article 7] requires the data controller to provide notice to the user of personal data transfer to third parties outside of the European Economic Area (EEA).

- Individuals could be informed of consent of disclosure, and given an **opportunity to withhold** it, if an employment, housing, education, healthcare, or life-altering decision is going to be made by an automated system. Similar considerations could be had for biometric information that is going to be analyzed, stored, or tracked. While there are many valid reasons why opting out could preclude access to employment, housing, education, healthcare, or other life-altering decisions, organizations should strive to minimize these instances except when strictly necessary—for example, as a means of security. In cases where a disability could prevent a person from participating, an opt-out or non-biometric alternative should always be provided, which should be able to result in provision of service.
- Each time data is sold or transferred to another party with separate governance, **consent for that transfer** must be obtained from the original data discloser.
- Platforms should strive to give users the ability to be granted, and **retain access** to, aspects of the platform even if users decline or retract consent critical to other portions.
  - There should be provisions for circumstances where users change their mind on consent—where either they offered consent and want to retract it, or they decide to grant further consent for more features.
  - In the public sector, where its users are granted rights, withholding consent for PII collection must not result in the denial of service. Provisions must be in place to provide rights afforded to residents and/or citizens regardless of data collection.

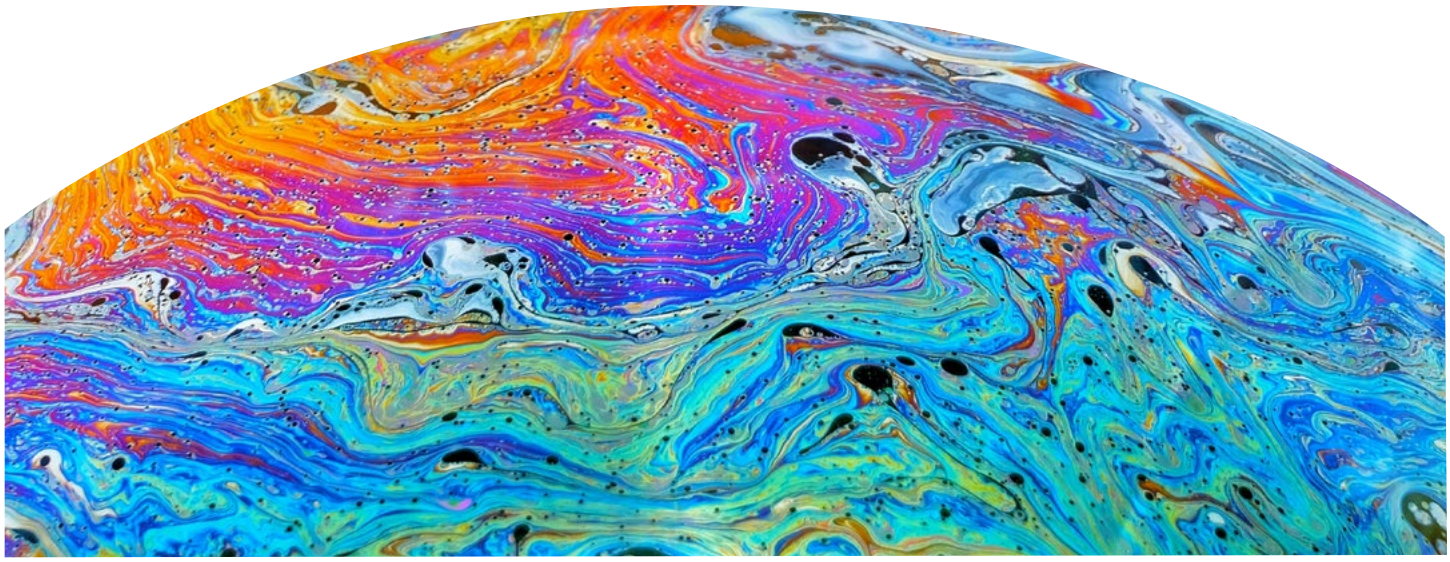

PROVENANCE

Policies must be sensitive to the fact that with the "big data" used in AI systems today, anonymity of individuals represented in the data is largely impossible for a sufficiently motivated actor.[6] Terms such as "pseudo-anonymous" are often reliant on contract terms or an assumption that the data won't be shared or combined with other data. As such, policies should strive to reflect the non-existence—or at least extreme difficulty—of anonymity in big-data and AI systems. Given this, combined with an increasing patchwork of global regulations, carrying the provenance of data throughout its lifecycle is critical to honoring rights guaranteed to individuals.

Verifiable audit > Retain agency > Protection of PII inferences > Source metadata

- **Metadata** can, and should, be leveraged to track and verify data sources and provenance.
  - All data about individuals should carry metadata that includes a robust (who/how/what/where/when) context of consent.
  - The provenance of a model (data subjects, purpose, governance, etc.) should be described through metadata that is amended throughout the model's lifecycle.
- Any inferences made regarding PII (age, gender, etc.) **should be treated as PII and protected as such**. In this context, inferred data is quasi-PII in that it is created with a prediction model, but if data was provided by a data subject, it would be covered as PII. Inferred data is a special type of synthetic data.

[Opinion No. 20-303, p. 15] California's attorney general has stated that "internally generated inferences that a business holds about a consumer are personal information within the meaning of the CCPA, and must be disclosed to the consumer on request." This is true even if the information on which the inferences are based was exempt from the CCPA when collected.

- Data disclosers should **retain agency** over how and for what purpose the data they disclose will be used.
  - Consent proxy delegation (human, algorithmic, or hybrid) follows data as it travels, allowing the discloser (or discloser's proxy) to keep meaningful control over the data as it changes hands and uses.
- Data disclosers should be able to **verifiably audit** how their disclosed data was used and for what that data was used, including editing/expunging PII or inferences made based on prior disclosures.

III. ARTIFICIAL INTELLIGENCE

AI, algorithmic, and autonomous systems are increasingly part of the products and processes of daily life. Omnipresence should not be mistaken for infallibility. In fact, any regulations for these systems should assume fallibility, check for it, include human review, and adapt to changing contexts—from development and deployment to maintenance.

Prioritize human consequence and agency > Reliance on AI

A human-centered approach is key to deciding where it is appropriate to apply AI.

Every algorithm, system, and model holds the possibility for error. Where insights derived from data could impact the human condition, the potential for harm at scale to individuals and communities should be the paramount consideration. Big data can produce compelling insights into populations, but those same insights can be used to unfairly limit an individual's possibilities in life. There are specific use cases for AI that require special consideration to mitigate the realization of severely adverse outcomes. Regardless of the governance approach—be it risk based, rights based, or principles based—use cases with even a slight potential for detrimental impacts, such as threats to health, personal freedom, or overall well-being, must undergo the strictest regulatory scrutiny, evaluating both whether and how they should be deployed.

Retrain (dynamic) models > Static models

Dynamic models preserve value and provide sustainability.

There needs to be consideration for how a model's data and decision-making ability will fare with time and shifting circumstances. Without retraining, a model is not just incomplete, but ineffective as a sustainable tool for the population it aims to serve. While models that can shift based on real-world context or discovered features should be pursued, dynamic models come with their own risks and may be more susceptible to bias feedback loops or adversarial attacks.[7]

Trustworthiness > Transparency

Transparency is a useful reform tool, but trust is what provides stability throughout an organization.

When it is genuine, transparency can be a critical component of an effective communications strategy, but an abundance of transparency about superfluous things can be used to distract from bigger, lesser-known issues. Being trustworthy is a higher bar than maximizing transparency for its own sake. To be trustworthy means supporting and advancing the values and interests of stakeholders and requires attending to establishing, building, maintaining, or repairing trust. This could manifest in many ways, and efforts to build trust in the public sector might look different than in the private sector. The point here is that the outcomes transparency tries to achieve are often better served through a focus on manifesting trust. The level of attainable transparency for any organization is capped by the opposing risk of too much transparency. Trust, on the other hand, is limitless.

Model an aggregate population > Model an individual

Avoid collecting excessive personal information, while deriving similar value with less risk.

Today's marketing "holy grail" is finding a way to communicate with an audience of one.[8] This, however, requires organizations to know a substantial amount about an individual, likely including PII. There are myriad risks involved in having such depth of information on so many people. Striving to minimize the amount of information needed while still achieving the same goal reduces risk and preserves value in the long term. In the case of marketing, all that is needed to achieve a comparable level of value is to be aware of which communications persona a stakeholder is most responsive to and to ensure that they are categorized in the right bucket at an appropriately relevant (yet minimal) level of specificity. This minimizes the amount of information needed for any single person and makes marketing operations much simpler—everyone wins.

DATA SCIENCE

Humans must be accountable for autonomous systems, and must hold agency over and meaningful control of those systems. All datasets and accompanying analytical tools carry a history of human decision-making. That history should be carried with the data for as long as possible, in order to remain auditable throughout its lifecycle.

Human intervention > Ethical review > Prohibition of deception > Auditable record

- An **auditable record** should be kept of the model itself, the algorithms and datasets used, the factors applied, the factors' corresponding weights, and the model outcomes. If changes are made to the data, algorithm, or model over time, the record should include, at a minimum, the actor, the change made, and the reason for the change.
  - In cases where an automated system makes a decision that affects a human's ability to proceed with their intended course of action, the metrics and reasoning of that decision should be transparent, understandable, accessible, and subject to recourse.
  - The benefits of an autonomous system should be evaluated based on its overall systemic impact, rather than on a single performance metric. For example, if the rise in availability of autonomous vehicles causes more use of car transportation overall, autonomous vehicles should be evaluated in terms of total auto fatalities, not just fatalities per vehicle mile traveled, and compared to manually driven cars.
  - Systems that self-optimize over time must be treated with an elevated risk profile. These systems should be routinely monitored for deviations and amplification of unwanted bias, and routinely subjected to a third-party, independent audit.
  - For generative AI systems based on large language models (LLMs) where the models are necessarily black boxes, it is paramount to retain model outputs for assessing drift, coherence, or any future audits that might be required.
  - Outputs from LLMs should be catalogued and watermarked with the ability for third parties to query for matching output.
- There should be a **prohibition of deceptive** or psychologically manipulative actions against humans by AI tools, whether material harms are experienced or not.

> The EU AI Act [Article 5] prohibits the use of an AI system that "deploys subliminal techniques beyond a person's consciousness in order to materially distort a person's behavior in a manner that causes or is likely to cause that person or another person physical or psychological harm."

  - AI systems should not be allowed to represent themselves as human agents.
  - As a default, AI-generated content should be labelled as such.
  - Autonomous systems seeking to mimic humans, whenever any party to an interaction is human, must disclose whether they are a machine, the responsible party, and the purpose of the interaction prior to gaining consent for the interaction from the human.

- Products and research practices should be subject to **ethical review** and/or third-party independent audit. Organizations should establish consistent, efficient, and actionable ethics review and audit practices for new products, services, and research programs.
  - Instate internal accountability measures that include representation of AI expertise on the board of directors. External ethics advisory boards can also be helpful. These governance bodies should be responsible for oversight of the implementation and continuous updating of monitoring and transparency measures.
- The need for human and machine collaboration cannot be understated. Determining the appropriate amount of **human intervention** is a question of balancing probability with consequence—a risk-based decision-making process.
  - When autonomous systems are being designed, developed, deployed, or repurposed for any new or novel contexts, an appropriately knowledgeable human—or governance body—should make go/no-go decisions for an initial period of time or number of instances. It might make sense to have this human oversight in perpetuity.
  - When considering or deploying autonomous systems for high-impact use cases (e.g., healthcare, safety), the necessity of a permanent human manager should be considered.

The EU AI Act takes a risk-based approach to regulating AI systems, in which the higher the risk of the implementation, the more restricted the use. Similarly, some private organizations take a risk-based approach with their internal policy, applying higher standards of governance and accountability to higher-risk AI uses.

In situations where an AI either acts as the safety component of a product or makes up the entirety of a product (including that product's safety component), the EU AI Act [Article 6] categorizes the AI as "high-risk," the most restricted category behind "unacceptable/prohibited." High-risk AI is subject to greater requirements and an ex-ante conformity assessment before use.

MACHINE LEARNING

The rapidly evolving nature of machine-learning (ML) applications necessitates organizational, and perhaps even national, interventions for rigorous testing, benchmarking, and auditing of ML-based systems for fairness, accuracy, and reliability across an appropriate and defined range of use contexts and populations.

Determined liability > Disclosure within transfers > Clear improvement > Attentively train

- Evaluate **training data** for ethical sourcing, and consistently monitor for new bias in source data.
  - For any data-ingest system for data scientists or AI developers, integrate tools or user experiences to help users understand and identify whatever relevant biases might be present in datasets that could be a risk to data quality.
- For any "black box" ML-based system to be deployed, make accessible information that states the improvements the ML system offers over the status quo and the ongoing protocol employed for measuring improvement on relevant metrics, so that the advantage of the AI system compared to the human/status quo option is **clear**.
- Each time a model is sold, shared, or transferred to another party with separate governance, this **transfer should be disclosed** in a way that includes: rules the implementer should follow; caveats about the data/algorithm/model; known biases the algorithm amplifies, or does not take into account (and if the model accounts for bias, disclosure regarding how it does so and its limitations); and what fairness issues might exist and what decisions have been made to mitigate fairness risk.
- When there is payment for models—as in models as a service (MaaS) business models—liability for the models' outcomes must be **determined** and documented as part of the transaction.

IV. PUBLIC SECTOR

Policymakers and public-sector organizations should consider how technology interfaces with and impacts broader communities, as they have a duty to ensure net societal benefits while protecting the public from harm. In the face of applying novel technology, they must balance the potential for profound benefit with minimizing disparate and negative impacts. They should also consider the opportunity to model governance behaviors and practices at the highest level.

Inclusive consideration > Utilitarianism

Protect and plan for the most vulnerable populations, who are often on the fringes of consideration.

In the face of potentially harmful impacts from technology, the public sector must prioritize the needs of the most vulnerable, in order to minimize the potential amplification of preexisting, discriminatory institutional structures. The Universal Declaration of Human Rights should act as a baseline standard whose provisions should be prioritized above all else. Where other sectors and contexts fail to consider certain populations due to minority status, disenfranchised identity, or other factors, the public sector must act as an advocate and a safety net. When considering the needs of the collective, these populations must be included in the whole. Rather than placing excessive weight on the experience and utility of the majority, governments must always weigh the risk of how the most vulnerable could be disproportionately affected.

Protection of the commons > Incentives for individuals

Consider the needs of the collective over the interests of individuals.

The "Tragedy of the Commons" describes a phenomenon in which a shared resource, from which no one can be excluded, is degraded over time due to each individual's incentive to get more out than they put in. Public organizations and services should strive, as much as possible, to protect, maintain, and bolster the public commons. In the context of technology's effects on society, the commons of public privacy have shifted. To avoid the detrimental effects of misaligned incentives, the public sector should prioritize the collective needs of the public and serve to set both guideposts and boundary lines for private behavior, preventing the private interests of individuals or organizations from infringing on the needs of the collective or any particular marginalized group. For example, a business that sells security cameras, like the individuals using them, is incentivized to install as many as possible, covering as wide an area as possible. While it is generally legal to record video and audio in public, the public sector often protects a "reasonable expectation of privacy," which includes stipulations around whether and where cameras can be placed, hidden, etc. These bounds should be informed by the values and priorities of the public, especially those most vulnerable, and should apply to the principles and functions of public organizations.

Empowered public > Informed public

Go beyond communication to collaboration with those affected by public decision-making.

Public participation is a value based on the belief that those who are affected by a decision have a right to be involved in the decision-making process—if they are not at the table, they are on the menu. It is a standard above an informed public because it requires seeking out, recognizing, and communicating the needs, interests, and ideas of those affected, and subsequently promises that the contribution of the participants will have influence on the decision. Often, this requires making the information and participation more broadly accessible.

Proactive iterations > Reactive incrementalism

Keeping pace with technology and its effects necessitates anticipation and creativity.

The pace of technological advancement is growing exponentially, and its impacts are too large and systemic to be approached with protocols designed for a previous decade's status quo. The public sector should lean into existing policy experimentation initiatives and expand their remit. Contemporary approaches to agile governance are focused on being responsive to stimuli, often taking the form of technological progress. What would be more impactful, for example, is to leverage data science for long-term policy. For instance, technology and data science give governments the capacity to understand hyper-local market conditions, as well as trends over time. Failing to tie these to relevant policies, as opposed to a one-size-fits-all approach, is a missed opportunity. For example, through Fannie Mae and Freddie Mac, the US government provides liquidity, stability, and affordability to the mortgage market by buying consumer loans from lenders. Twenty years ago, Fannie Mae and Freddie Mac had one maximum loan amount for the continental United States, and another for Alaska and Hawaii, arguing that home prices in those states were structurally different from those in the rest of the country. When markets such as New York and San Francisco became more expensive than Alaska and Hawaii, the failure to be more proactive to local markets highlighted dysfunction in housing policy. Today, the United States is still a long way from locally indexed housing policy, but Fannie and Freddie now treat a few dozen counties in the continental United States as similar to Alaska and Hawaii.[9] Governance bodies should be leveraging these capabilities to enshrine new policies that proactively iterate in dynamic and responsive ways, where interventions can still happen but are the exception, rather than the rule.

AUTONOMOUS SYSTEMS

Public-sector policies around autonomous systems should reflect the priority of just and quality service to society. Public-sector organizations, in particular, should consider investment in the veracity and transparency of their models.

National-level safety standards > Empowering sectors > Addressing disparate impact > Public performance disclosure

- For any government use of machine learning, neural networks, or other AI techniques, there should be **public disclosure** of performance metrics on bias, fairness, accuracy, and reliability, as well as consumer notice of associated risks and limitations. These metrics should always include performance of the AI-based system compared to performance of the status quo.
- For any government use of machine learning, neural networks, or other AI techniques, there should be mandatory funding of regular (quarterly, biquarterly) studies that **report on disparate impacts** of protected classes, ideally conducted by a neutral third party. Testing protocols must be published.
  - If statistically relevant disparities are discovered, the system should be retuned within a prompt, predefined time period (e.g., thirty days), and repeated until the problem is solved. If the system is producing disparate impact for more than a predefined limit (e.g., three consecutive months), its use must cease, and humans should take over until the system can remove disparate impacts that are more significant than human-administered policies.
- Governments should **empower sector-specific agencies** (like health, education, criminal justice, and welfare) to audit and monitor high-risk autonomous systems within their domains to reflect their own histories, regulatory frameworks, and hazards.
- There is a need for a **national AI safety body** to determine a risk-, rights-, or principles-based approach to regulating AI systems. It would be tasked with review of proposed autonomous systems in addition to governance of existing ones. Similar to how other standards bodies work today (e.g., the National Transportation Safety Board), the safety body sets a minimum bar and sectoral agencies (e.g., the Environmental Protection Agency) can set higher standards.

ACCOUNTABILITY AND RECOURSE

Recourse is fundamental to democracy and human rights. To be at all compatible with that given, autonomous systems applied in the public sector must be held to the highest standards of recourse, as any public service would be. Policy and legal experts should have regular, open conversations with engineers and technologists to further their understanding of, and adaptability to, new technologies in order to fully understand that government has the highest level of responsibility to be held accountable for their use.

Appeal and remedy > Robust decision-making > Audits and understanding

- Vendors and developers who create AI and automated decision systems for use in government should, short of a court order to the contrary, be made to waive any trade secrecy or other legal claim that inhibits full **auditing and understanding** of their software.
- On any "black box" system in the public sector that impacts the safety and freedom of humans (no-fly lists, parole, access to government entitlements, health-insurance risk pooling, etc.), the models applied to such high-impact decisions should be made based on a **robust variety of indicators**, be routinely tested for disparate impacts, and have results of these studies published. Furthermore, investments should be made in improving these systems by making them explainable.
- There must be a trackable and auditable means of **appeal and remedy** when access or services are unfairly denied or delayed by algorithmic or AI error (e.g., flight registries, access to public entitlements). The remedy must happen within an established time period, and measures should be taken to ensure improved fairness for others with similar use cases.
  - Individual "subjects" of the system must be able to gain an explanation of why they were selected, given an immediate opportunity to validate or refute their selection, and—in the absence of sufficient evidence from the entity in power or ability to explain the "selection"—must be given immediate redress of the situation. In cases where facts are contrary to system results, that redress should be permanent.

DATA USE AND SECURITY

Governments must hold themselves to the highest standards of data security and, in addition, set the outer limits for what is deemed an appropriate use of data at large. With the understanding that markets are incentivized to use the highest-value data for its highest-value use, there must be guardrails and restrictions to prevent possible infringements on peoples' lives, liberties, or pursuits of happiness. Data has enormous potential to benefit the public sector, specifically regarding the creation and upkeep of legislation and other legal measures. Data coming from or used in key public-sector industries like healthcare, transportation, law enforcement, or judicial systems must be given especially strict consideration.

Safety net for private conduct > Biometric data restrictions > Data-informed regulation > Protection of life and freedom

- Restrictions and bans need to be considered for high-risk uses of data where errors may cause loss of life or freedom (sentencing and parole decisions, no-fly lists, medical risk assessment). The **protection of free, quality life** must be treated as the utmost priority.
- Data, in robust and bias-checked form, should be used to **inform regulation**. This gives legislators a more intimate relationship with data and a deeper understanding of its strengths and weaknesses.
  - Legislation of data-indexed events should strive to be dynamic—responsive to real-world changes—rather than static (e.g., retirement age based on life expectancy of a community or conforming loan limits based on area median prices rather than a universal, one-size-fits-all number).
- No **biometric systems**, including genomics mapping and facial and affect recognition, should ever be used as a barrier to access for public services, employment, housing, education, healthcare, etc.
- No continuously aggregating databases of **biometric information** should be held by any entity other than law enforcement.
  - If law enforcement maintains a biometric database, that information cannot be connected to the internet in any way. Every access to that system of record must be premised by a warrant and be auditable through a Freedom of Information Act (FOIA) request or petition to a court. Access to biometric information must be rigorously prevented from being used for discriminatory "predictive" practices.
  - Prohibition of DNA databases held by law enforcement should be considered. At the very least, these databases must be subject to highly stringent restrictions against being aggregated, licensed, or sold. Prohibitions must also require that access be exclusively for known individuals, and access logs must be routinely vetted and tested for foul play.

The EU AI Act [5.2.3] explicitly lists "biometric identification and categorization of natural persons" as a high-risk use of AI. It proposes [Article 43] specific restrictions and safeguards to the uses of remote biometric identification systems for the purpose of law enforcement.

  - DNA databases should not be used to search and contact groups of people. It should remain acceptable for law enforcement, with a detailed and specific court order, to identify specific individuals for a specific reason.
  - DNA should not be used for feature extraction or the creation of synthetic data that could then be used to inform and/or train an AI process or model.

- Governments must fulfill their role as a **safety net for ethical conduct** in the private sector.
  - Under the purview of antitrust, platform- or ecosystem-based companies should be restricted from competing with others on their platform by using data derived from that company's participation in the ecosystem or platform. For instance, if an ecommerce ecosystem company learns that there are great profits to be made with linen products, and it learned this from linen vendors on its platform, the platform company should not be allowed to enter this market.
  - Clear guidelines need to be established about data collected for medical research and the commercialization of that data—whether or not it is "anonymized."
  - Prohibition of tracking databases (e.g., automatic license-plate reading and cataloging, smartphone location data) built and/or maintained by private companies should be considered. At the very least, there must be highly stringent controls over these "dragnets" to prevent data from being aggregated, licensed, or sold. Prohibitions must also require access to be exclusively from known individuals, while access logs must be routinely vetted and tested for foul play.

V. GOVERNANCE

As technology becomes as fundamental to the functioning of an organization as its board of directors and employees, a fundamental shift is needed in the way responsibility and accountability are distributed.

Being a responsible body—whether that means a development team, an entire organization, or a nation-state—now includes accountabilities for all the inputs, outputs, impacts, hidden costs, and externalities of the technology tools in purview. The only way to achieve the level of insight needed is to develop a culture in which governance is so embedded and routine that it is second nature, and in which engaging with governance is commonplace. This exists today in regulated industries such as financial services, but less regulated industries can, and should, exercise this muscle too.

Minimize harm (to stakeholders) > Maximize value (for shareholders)

Risk mitigation and harm minimization—while maximizing benefits for all—are essential to any long-term value strategy.

Above all else, technologies should respect the persons subjected to them, particularly when they are used implicitly or without specific consent. When technologies are used to unfairly limit an individual's possibilities, meaningful harm occurs. At scale, these harms can be as atrocious as genocide. The issue is serious. Even minor harms can compound and scale, creating broad disadvantages for inadvertently targeted segments of the population.

No money is worth that societal cost. If an organization values its stakeholders above shareholders, then the choice to minimize harm to individuals over maximizing short-term revenue is always the right choice. In the long term, technology can maximize its value to all by being averse to harm.

Value stays with data subject/discloser > Data collector/aggregator/user

Ensure a robust data ecosystem to maximize the value that stays with data disclosers.

When data providers retain more value from the information they share, they are more likely to continue sharing data. If all the value resides with the data collector, the incentives for more data disclosure begin to deteriorate. To sustain a thriving data ecosystem, it is crucial to guarantee that data providers retain a significant portion of this value. This approach fosters a generative environment for data-driven ecosystems, offering increased opportunities for innovation among data collectors, aggregators, and end users, and ultimately benefiting the public.

Fairness through "values transparency" > Enforcing equality

Focus on creating a level playing field and disclose the values that drive that decision-making.

Equality is when everyone gets the same thing, regardless of their personal needs or situation. Equity happens when people are given what they need in order to engage fairly with others. When aspiring to equity or fairness, it is important to note that the correct approach is largely subjective, as it depends on an interpretation of need. In the context of AI, fairness is in demand. But the only way to truly understand how an organization is optimizing for its unique definition of fairness within its autonomous systems is to understand the values for which it is prioritizing and optimizing.

Manage internalities > Externalize internalities

Minimize potential harms with robust internal governance, and before harms have a chance to scale.

The greatest advantage—and greatest risk—of digital technologies is their ability to scale. Relatively small oversights in AI governance can lead to radically outsized harm to communities and existential risks to the organizations that proliferate them. Having robust internal governance practices go a long way toward minimizing this risk, but it is still necessary to have a plan of accountability in place for when unintended harm occurs.

RECOURSE AND REDRESS

Efficient and effective methods of recourse and transparent processes for redress are essential to the transparency, accountability, and sustainable functioning of any system, especially in cases in which autonomous systems are carrying out decisions that would otherwise have been executed by a human and/or are affecting the decisions of other actors. At the very least, autonomous systems should be subject to the same recourse and redress obligations as a person would be in a comparable situation. Additional measures should be heavily considered to account for the scale of impact possible with autonomous systems and the inability of these systems to self-monitor and self-correct.

Reporting of redress > Ease of recourse > Clear recourse > Protect compliance

- **Provide protections** for conscientious objectors and ethical whistleblowers.
  - Ethical malfeasance of data practices for autonomous systems should be reported within companies.
  - One way to encourage disclosure and public reporting is for an external governance body (e.g., the US Securities and Exchange Commission) to provide a ceiling for liability in exchange for disclosure.
- All autonomous systems must have **clear methods of recourse.** Central to this is a clearly identified way that users or those subjected to the technology can escalate concerns to responsible engineers, internal governance, and external governance. There must be clear service commitments (i.e., redress) for accepting, tracking, and reporting these escalations to entities with adequate authority and capability, which could be internal and/or external.
  - To build institutional knowledge, recourse requests can be used as a valuable source of feedback. Regular reporting to product-management and executive teams can help to ensure the value of this feedback loop can be maximized. Public disclosure of these reports is something leading companies will disclose in corporate social responsibility (CSR) AI transparency reports.
- On any autonomous system, there must be a simple, **clear, omnipresent method for a human to notify** system designers or engineers of anomalous system behavior. Ideally, this will happen in a single step as opposed to a lengthy, difficult-to-find menu option. This method should be accessible via any touchpoint in the system, trackable by the person making the report, and auditable by an internal or external governance body.
  - In limited cases where the anomaly is serious in scale and it is universally safe to do so, this single step could also shut down the system.
- **Redress** should happen when an organization receives a request for recourse. An organization should have clear policies, standards, and commitments for when that request for recourse is acted upon, by which role, who is in that role, and what the expected outcomes are. Leading organizations will also choose to report on their own redress performance to stakeholders and the public.

DATA AND SECURITY

The bar for the handling of data should be to meet or exceed the expectations of data disclosers. This applies to data-aggregation methods, cybersecurity measures, distribution, and disposal. To use the resource of data, and leverage its immense potential, the policies put in place to ensure it is handled responsibly must be thoroughly considered.

New data-ownership model > Protected class of data applications > Integrated security

- **Security measures should be integrated** with the concept of data use, and should necessarily be included as a part of any data product or service.
  - When providing tools or services that allow average individuals to input, collect, and/or exchange data, resources—such as a set of standard practices or settings—1should be included so security is a feature everyone can enjoy.
- Certain types of data or certain uses of autonomous systems need to be given special consideration. As the only global set of ratified human-rights protections, the Universal Declaration of Human Rights is a strong foundation—not only because the declaration forms part of customary international law, but because it offers a way to apply moral and diplomatic pressure to governments that violate any of its articles. The declaration should be the foundation for zero-tolerance policies that set the protective outer boundaries for further policy and ethical deliberation, while ensuring the most essential of human rights are never breached. What's more, the organizations that offer AI systems may want to explicitly offer a set of rights to their stakeholders.
  - Protecting marginalized populations from misinformation is particularly important as technology continues to pose significant risk to human rights at scale.[10] When this is allowed to occur, or safeguards are insufficient to prevent these abuses from happening, the consequences can be pervasive. This misinformation can be as simple as false information proliferated by trusted parties or governments or as complex as deep fakes, which are algorithmically manipulated or created video with the intention to mislead. The consequences for proliferating this harmful content must have meaningful enough consequences to sufficiently discourage harmful behavior. If harm can scale exponentially, so must the criminal penalties, including those for the responsible parties in a position to prevent such outsized risks (including those that offer the tools).
  - There should be a ban on the sale of financial, location, DNA, and health data. These datasets represent surveillance capacity that should only be available to governments through a court order.
  - Local, community-based governance bodies should enjoy the greatest autonomy in accepting or rejecting the use of any technologies, in both public and private contexts, which could limit human rights or the fundamental freedoms they are meant to protect.

The EU AI Act [Article 7] states that, beyond the explicitly listed unacceptable and high-risk uses of AI, an AI system may be added as high risk if it meets some additional conditions including: posing a risk of or causing harm to health and safety; posing a risk of or causing adverse impact on fundamental rights; potentially harming or adversely impacting persons who depend on the outcome produced with an AI system and/or cannot reasonably opt out; potentially harming or adversely impacting users who are in a vulnerable position, in particular due to an imbalance of power, knowledge, economic or social circumstances, etc.

- Incorporate a **new model for data ownership** in which ownership resides with the subject the data describes. Within this model, the only thing a corporation or public entity can claim is a nonexclusive license to use that data. That license should then be revocable at any time for any reason by the data subject. A few Web 3.0 technologies are offering these capabilities today.

CONCLUSION

An approach that acknowledges and outlines the value tradeoffs always present in decision-making enables the intentionality and innovation necessary for public and private policymaking to keep up with the current pace of technological change. Leaders in the public and private sectors have an opportunity to use this approach to push standards of responsibility in technology beyond the current regulatory standard. This framework supports the technology sector's collective effort to raise the bar to which others will aspire, standardize best practices to positive impact, and improve outcomes for all individuals. Will you be a part of setting the new bar or playing catch-up?

ENDNOTES

1 "A European Approach to Artificial Intelligence," European Commission, June 2023, https://digital-strategy.ec.europa.eu/en/policies/european-approach-artificial-intelligence; "Regulatory Framework Proposal on Artificial Intelligence," European Commission, last visited August 15, 2023, https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai; "AI Risk Management Framework," National Institute of Standards and Technology, March 30, 2023, https://www.nist.gov/itl/ai-risk-management-framework; "Blueprint for an AI Bill of Rights," White House, March 16, 2023, https://www.whitehouse.gov/ostp/ai-bill-of-rights/; "FACT SHEET: Biden-Harris Administration Announces Key Actions to Advance Tech Accountability and Protect the Rights of the American Public," White House, October 4, 2022, https://www.whitehouse.gov/ostp/news-updates/2022/10/04/fact-sheet-biden-harris-administration-announces-key-actions-to-advance-tech-accountability-and-protect-the-rights-of-the-american-public/; "Artificial Intelligence," Organisation for Economic Co-operation and Development, last visited August 15, 2023, https://www.oecd.org/digital/artificial-intelligence/.

2 "Privacy," Business Roundtable, last visited August 15, 2023, https://www.businessroundtable.org/policy-perspectives/technology/privacy.

3 Steven Tiell and Lara Pesce Ares, *Principles to Practice: Using Ethical Spectrums to Guide Decision-Making, Atlantic Council*, July 28, 2022, https://www.atlanticcouncil.org/in-depth-research-reports/issue-brief/principles-to-practice-using-ethical-spectrums-to-guide-decision-making/.

4 For more on leveraging values to operationalize ethics, see: John Basl, Ronald Sandler, and Steven Tiell, *Getting from Commitment to Content in AI and Data Ethics: Justice and Explainability, Atlantic Council*, August 5, 2021, https://www.atlanticcouncil.org/in-depth-research-reports/report/specifying-normative-content; "Operationalising Ethics and Accountability Workbook," Monetary Authority of Singapore, 2022, https://www.mas.gov.sg/-/media/MAS-Media-Library/news/media-releases/2022/Veritas-Document-3B---FEAT-Ethics-and-Accountability-Principles-Assessment-Methodology.pdf#page=24.

5 Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (New York: Crown, 2016).

6 Scott Berinato, "There's No Such Thing as Anonymous Data," *Harvard Business Review*, February 9, 2015, https://hbr.org/2015/02/theres-no-such-thing-as-anonymous-data; Yves-Alexandre de Montjove, et al., "Unique in the Shopping Mall: On the Reidentifiability of Credit Card Metadata," Science 347, 6221 (2015), http://www.sciencemag.org/content/347/6221/536.full; Gregory S. Nelson, "Practical Implications of Sharing Data: A Primer on Data Privacy, Anonymization, and De-Identification," SAS, 2015, http://support.sas.com/resources/papers/proceedings15/1884-2015.pdf; Natasha Singer, "With a Few Bits of Data, Researchers Identify 'Anonymous' People," *New York Times*, January 29, 2015, http://bits.blogs.nytimes.com/2015/01/29/with-a-few-bits-of-data-researchers-identify-anonymous-people/?_r=0; Abdul Majeed and Sungchang Lee, "Anonymization Techniques for Privacy Preserving Data Publishing: A Comprehensive Survey," IEEE Access 9 (2021), 8512-8545, https://ieeexplore.ieee.org/document/9298747.

7 "Features" in AI systems are individual independent variables or attributes that are used to make predictions or infer patterns. In a credit-approval model, features might be an applicant's annual income, employment status, credit score, credit history, total outstanding debt, and many other datapoints. Each feature has an associated "weight" or influence over the output of a model.

8 An "audience of one" refers to the idea of tailoring marketing efforts and communications to resonate with individual consumers as if they were the sole recipient. This approach moves away from broad, generalized marketing campaigns to hyper-personalized strategies that appeal to specific needs, preferences, behaviors, and circumstances of each individual and, as such, require tremendous amounts of data, often including PII.

9 "Conforming Loan Limit Values Map," Federal Housing Finance Agency, last visited August 15, 2023, https://www.fhfa.gov/DataTools/Tools/Pages/Conforming-Loan-Limit-Map.aspx.

10 Ruth Hickin, "How Is Technology Affecting Our Human Rights?" World Economic Forum, December 11, 2017, https://www.weforum.org/agenda/2017/12/how-are-today-s-biggest-tech-trends-affecting-human-rights/; "Urgent Action Needed over Artificial Intelligence Risks to Human Rights," UN News, September 15, 2021, https://news.un.org/en/story/2021/09/1099972.

CONTRIBUTORS

The specific policies and ideas contained herein were developed in collaboration with several key contributors:

STEVEN TIELL is a nonresident senior fellow with the Atlantic Council's GeoTech Center. He is a strategy executive with wide technology expertise and particular depth in data ethics and responsible innovation for artificial intelligence. Tiell has helped to build de facto public policy by helping the Monetary Authority of Singapore to write new regulatory guidance for fairness, ethics, accountability, and transparency. Since embarking on data ethics research in 2013, Tiell has contributed to and published more than a dozen papers and has consulted with dozens of organizations in advertising, defense, financial services, government, high-tech, media, public policy, public safety, and telecom industries. He often speaks on topics such as governance, trust, data ethics, synthetic media, misinformation and disinformation, and industry trends. Much of Tiell's career has focused on leveraging technologies for impact in the public sector. In addition to his fellowship at the Atlantic Council, he is the vice president for strategy at DataStax, a Unicorn database-as-a-service company with open-source technology and mindset at its core. Prior to DataStax, Tiell founded and led Accenture's global data ethics and responsible innovation practice. While at Cisco, he led global programs for Connected Urban Development and went on to lead the Smart+Connected Communities Institute. He has also led product research and development for a digital media firm and started several entrepreneurial ventures. Earlier in his career, he designed and built public safety systems at Northrop Grumman and helped to operate one of the nation's first nonprofit internet service providers. Tiell lives in San Francisco with his family and holds a BS in computer science engineering and information systems and an MBA in sustainable management."

LARA PESCE ARES is a Responsible Innovation consultant at Accenture focusing on the design, development and implementation of emerging technologies in a responsible way. Passionate about the power of scaled impacts, she collaborates frequently with private and public organizations at the forefront of technological and social innovation, focusing on key themes of bias, inclusion, and human-centered design. She was previously an emerging tech trends author of the Accenture Technology Vision, where she authored pioneering trends on technology democratization and Web 3. She holds a BA in public policy from New York University, where she did a senior thesis on the landscape of data-driven initiatives in city governments."

BEAU CRONIN
WeaveGrid, previously the Data Guild

DAVID DANKS
University of California, San Diego, previously Carnegie Mellon University

Atlantic Council

Board of Directors

CHAIRMAN
*John F.W. Rogers

EXECUTIVE CHAIRMAN EMERITUS
*James L. Jones

PRESIDENT AND CEO
*Frederick Kempe

EXECUTIVE VICE CHAIRS
*Adrienne Arsht
*Stephen J. Hadley

VICE CHAIRS
*Robert J. Abernethy
*Alexander V. Mirtchev

TREASURER
*George Lund

DIRECTORS
Stephen Achilles
Elliot Ackerman
*Gina F. Adams
Timothy D. Adams
*Michael Andersson
Alain Bejjani
Colleen Bell
Sarah E. Beshar
Stephen Biegun
Linden P. Blue
Brad Bondi
John Bonsell
Philip M. Breedlove
David L. Caplan
Samantha A. Carl-Yoder
*Teresa Carlson
*James E. Cartwright
John E. Chapoton
Ahmed Charai
Melanie Chen
Michael Chertoff
*George Chopivsky
Wesley K. Clark
*Helima Croft
Ankit N. Desai
Dario Deste
Lawrence Di Rita
*Paula J. Dobriansky
Joseph F. Dunford, Jr.
Richard Edelman
Stuart E. Eizenstat
Mark T. Esper
Christopher W.K. Fetzer
*Michael Fisch
Alan H. Fleischmann
Jendayi E. Frazer
*Meg Gentle
Thomas H. Glocer
John B. Goodman
Sherri W. Goodman
Marcel Grisnigt
Jarosław Grzesiak
Murathan Günal
Michael V. Hayden
Tim Holt
*Karl V. Hopkins
Kay Bailey Hutchison
Ian Ihnatowycz
Mark Isakowitz
Wolfgang F. Ischinger
Deborah Lee James
*Joia M. Johnson
*Safi Kalo
Andre Kelleners
Brian L. Kelly
John E. Klein
*C. Jeffrey Knittel
Joseph Konzelmann
Keith J. Krach
Franklin D. Kramer
Laura Lane
Almar Latour
Yann Le Pallec
Jan M. Lodal
Douglas Lute
Jane Holl Lute
William J. Lynn
Mark Machin
Marco Margheri
Michael Margolis
Chris Marlin
William Marron
Gerardo Mato
Erin McGrain
John M. McHugh
*Judith A. Miller
Dariusz Mioduski
*Richard Morningstar
Georgette Mosbacher
Majida Mourad
Virginia A. Mulberger
Mary Claire Murphy
Julia Nesheiwat
Edward J. Newberry
Franco Nuschese
Joseph S. Nye
*Ahmet M. Ören
Ana I. Palacio
*Kostas Pantazopoulos
Alan Pellegrini
David H. Petraeus
*Lisa Pollina
Daniel B. Poneman
*Dina H. Powell McCormick
Michael Punke
Ashraf Qazi
Thomas J. Ridge
Gary Rieschel
Charles O. Rossotti
Harry Sachinis
C. Michael Scaparrotti
Ivan A. Schlager
Rajiv Shah
Wendy R. Sherman
Gregg Sherrill
Jeff Shockey
Ali Jehangir Siddiqui
Kris Singh
Varun Sivaram
Walter Slocombe
Christopher Smith
Clifford M. Sobel
Michael S. Steele
Richard J.A. Steele
Mary Streett
Nader Tavakoli
*Gil Tenzer
*Frances F. Townsend
Clyde C. Tuggle
Francesco G. Valente
Melanne Verveer
Tyson Voelkel
Michael F. Walsh
Ronald Weiser
*Al Williams
Ben Wilson
Maciej Witucki
Neal S. Wolin
Tod D. Wolters
*Jenny Wood
Guang Yang
Mary C. Yates
Dov S. Zakheim

HONORARY DIRECTORS
James A. Baker, III
Robert M. Gates
James N. Mattis
Michael G. Mullen
Leon E. Panetta
William J. Perry
Condoleezza Rice
Horst Teltschik
William H. Webster

**Executive Committee Members*

List as of January 1, 2024

Atlantic Council

The Atlantic Council is a nonpartisan organization that promotes constructive US leadership and engagement in international affairs based on the central role of the Atlantic community in meeting today's global challenges.

Atlantic Council

1030 15th Street, NW, 12th Floor,
Washington, DC 20005

(202) 463-7226, www.AtlanticCouncil.org