Principles to Practice: Using Ethical Spectrums to Guide Decision-Making

Steven Tiell
Lara Pesce Ares

Atlantic Council
GEOTECH CENTER

Cover: Daniele Levis Pelusi/Unsplash

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

For more information, please visit
www.AtlanticCouncil.org.

July 2022

Principles to Practice: Using Ethical Spectrums to Guide Decision-Making

Steven Tiell
Lara Pesce Ares

TABLE OF CONTENTS

I.   INTRODUCTION

Societies are increasingly being shaped by technological change, and the pace is increasing exponentially. Every day, organizations make decisions that participate and shape this global transformation. As new technologies unlock unprecedented capabilities—and do so at scale—they also have the potential to bring about unprecedented existential risk. Organizations are being defined by their ability to manage these risks with a global perspective in mind, because the impacts of these decisions—intentional or not, direct or indirect—shape an organization's role in the ongoing, global, digital transformation, often with societal implications. And the intentionality with which an organization handles decision-making in this new era will be a differentiating factor in the marketplace.

As expectations shift rapidly beneath their feet, organizations have adopted a diverse set of strategies to manage these new risks. And, leaders know it is insufficient to inform decisions through legal bounds alone. Regulatory bodies in government struggle to keep up with the pace of digital change and have, thus far, failed to demonstrate a consideration of future risks in policies intended to be forward looking. As a result, existing laws and regulations easily become outdated, ineffective, and often mis-calibrated to current threats. In this context, even perfect compliance means potential exposure to existential risks—and when these are risks to the fabric of societies, action must be taken for the benefit of all. This means that to lead in this space, it's insufficient to follow any existing compliance framework. Leaders must set new ones.

In this proposed framework, the overarching bias is to protect the continued sustainability and enrichment of the human condition. To aid in this endeavor, this paper adopts a useful framework from the 2001 Manifesto for Agile Software Development—the notion that technology practitioners guide their decision-making by applying values. These values are constructed in context-sensitive ways, with the understanding that the pathway from values, through principles, and to action is critical, making this a unique approach.

Throughout this paper, there will be "value spectrums" where one idea is valued *over* another. This is denoted as a greater-than sign, ">". In practice, the spectrums are contextually relevant; here, they are used as examples for discussion. Designers, product managers, and development teams use these types of spectrums as guides for ethical decision-making. They do not prescribe any one correct answer; where any decision eventually lands on the spectrum is less important than the potentially stakeholder-rich deliberation that supports the final decision. Wherever an organization chooses to land along any one of the spectrums, the deliberate process of evaluating ethical priorities will necessarily be informed by the organization's values. In this manner, the organization is empowered to draw clear through-lines from core values to the features in their products and services—and, ultimately, to their communications, facilitating an intentional and trusted relationship with their customers, users, and the public. This approach serves to curate a deliberate and informed company culture, and further serves to protect digital companies from the existential risks their own decisions could foster.

> **"Where any decision eventually lands on the spectrum is less important than the potentially stakeholder-rich deliberation that supports the final decision."**

The overall goal of this framework is to recognize and respect the role that technology plays in the advancement of societies, while also recognizing the collective interest of societies to ensure the safety and security of individuals and groups. These spectrums are intended as a guiding tool to aid organizations attempting to walk a fine line between continuing to embrace the advancement of technology and realizing economic prosperity, without compromising their own values or their accountability to society.

II. GOVERNANCE

As technology becomes as fundamental to the functioning of organizations as their boards of directors and employees, there needs to be a fundamental shift in the way responsibility and accountability are distributed.

Whether it's a development team, an entire organization, or a nation-state, being a responsible body 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. Some spectrums to serve as a starting point might include the following:

Minimize harm > Maximize value

Risk mitigation and harm minimization are essential to any long-term value strategy.

As the Business Roundtable's Statement on the Purpose of a Corporation advocates, "companies should be led for the benefit of all stakeholders—customers, employees, suppliers, communities, and shareholders."[1] Above all else, technologies should respect the persons subjected to them, particularly when used covertly or without consent. When technologies are used to unfairly limit an individual's possibilities, meaningful harm occurs. When this happens at scale, genocide can occur. It's serious. While no decision can perfectly account for all possibilities, every reasonable effort should be made. Even ripples from small slights, at scale, can have harmonic amplification, creating tidal waves of disadvantage for inadvertently targeted segments of the population.

No money is worth that societal cost. And, if a company values its stakeholders above shareholders, then the choice to minimize harm to individuals over maximizing (short-term) revenue is always the right choice.

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

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

The more value retained by those providing data, the more apt they are to continue providing data. If all of the value resides with the data collector, the incentive structures for more data disclosure begin to deteriorate. To maintain a robust data ecosystem, it's important to ensure data disclosers retain a substantial amount of value. This breeds a generative environment for data-centric ecosystems that is in abundance, giving more opportunities for innovation to the data collectors and aggregators and, ultimately, users.

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, regardless of their needs or situation. Equity happens when people are given what they need to engage fairly with others. With artificial intelligence (AI), "fairness" is in demand, and the only way to understand how an organization is optimizing for its unique definition of fairness is to understand the values it cares about, and which it is prioritizing and optimizing.

Manage internalities > Externalize internalities

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

The greatest advantage of digital technologies is their ability to scale. Similar to cataclysmic environmental harms from bad industrial actors (e.g., rivers catching fire, Chernobyl's meltdown), relatively small oversights in AI governance can lead to radically outsized harms to communities, and existential risk to the organizations that proliferate them. Having robust internal governance practices go far to minimize this risk, but it is still necessary to have a plan of accountability in place for when unintended harms occur.

---

1   "Business Roundtable Redefines the Purpose of a Corporation to Promote 'An Economy That Serves All Americans'," Business Roundtable, August 19, 2019, https://www.businessroundtable.org/business-roundtable-redefines-the-purpose-of-a-corporation-to-promote-an-economy-that-serves-all-americans.

III. DATA PROCUREMENT AND USE

The amount of investment an organization needs to place in minimizing harm is directly related to the amount of value it derives from digital products or services that are informed by data. The value of data is increasingly compromised if the methods of procurement or use fall short of local laws or stakeholder expectations.

Being thoughtful across architecture, product development, and policy design doesn't just protect organizations by mitigating risk; it can also generate new value by improving relationships and retention with existing stakeholders, as well as attracting new ones.

These are samples of spectrums an organization might use to make data procurement and use decisions.

Collect relevant data > Any/everything possible

Miminizing data collection leads to better analysis and less risk.

It's always best to first consider the questions for which the answers could, and should, be informed by data. After the set of questions are articulated, data maps can be created to specify the data that needs to be collected. Then, data scientists can consider data-minimization techniques to further reduce the data needed to answer the questions. Doing so minimizing the data burden—the infrastructure, processes, and personnel required to handle large volumes of data—leaving the organization in a strong strategic position with minimal data risk should a breach or leakage happen.

Informed consensual use of data > Exploratory use

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

The more specific and informed the consent-sourcing process, 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 their data 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 more security and privacy risks increase—and all the while, value and public perception are degraded. All data has a useful life. Leadership and design teams should consider this as part of security protocols, consent regimes, and policymaking.

IV. ARTIFICIAL INTELLIGENCE

As the powerful tools of autonomous systems and artificial intelligence continue to define the products and processes of daily life, it is imperative to regulate them as the fallible tools that they are, and to ensure that at every stage—from development and deployment to maintenance—humans are at the center.

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 certain specific use cases for AI that require special consideration to mitigate the realization of severe adverse outcomes. Given the severity of consequences, such as risks to public health and safety, or even the loss of personal freedom, it may be appropriate to allow for appropriate governance methods that address fundamental AI deployment questions.

Re-train (dynamic) models > Static models

Dynamic models preserve value and provide sustainability.

There needs to be consideration of 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 valuable and sustainable tool for the people it aims to serve.

Be trustworthy > Transparent

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

When it's genuine, transparency can be a critical component of an effective communications strategy, but it can also be used to distract. Being trustworthy is a higher calling. To be trustworthy means attending to establishing, building, maintaining, or repairing trust at every opportunity and through many avenues. It could be answering the phone, immediately, without long, microtargeted phone trees. Or, it could be having customer service and sales agents trained on how to respond to end-user privacy concerns. Trust manifests in myriad ways; seizing the maximum number of opportunities to reinforce trust is a strong strategy to avoid unnecessary risk.

Model an aggregate population > Model an individual

Practice "clustering" to avoid excess collection of personal information; aim to derive similar value with less risk.

Today's marketing holy grail is to communicate with an audience of one, but this requires organizations to know a substantial amount about an individual, likely including personally identifiable information (PII). There are myriad risks involved in having such depth of information about so many people. The use of clustering can minimize the amount of information needed for any single person and make marketing operations much simpler, so everyone wins.
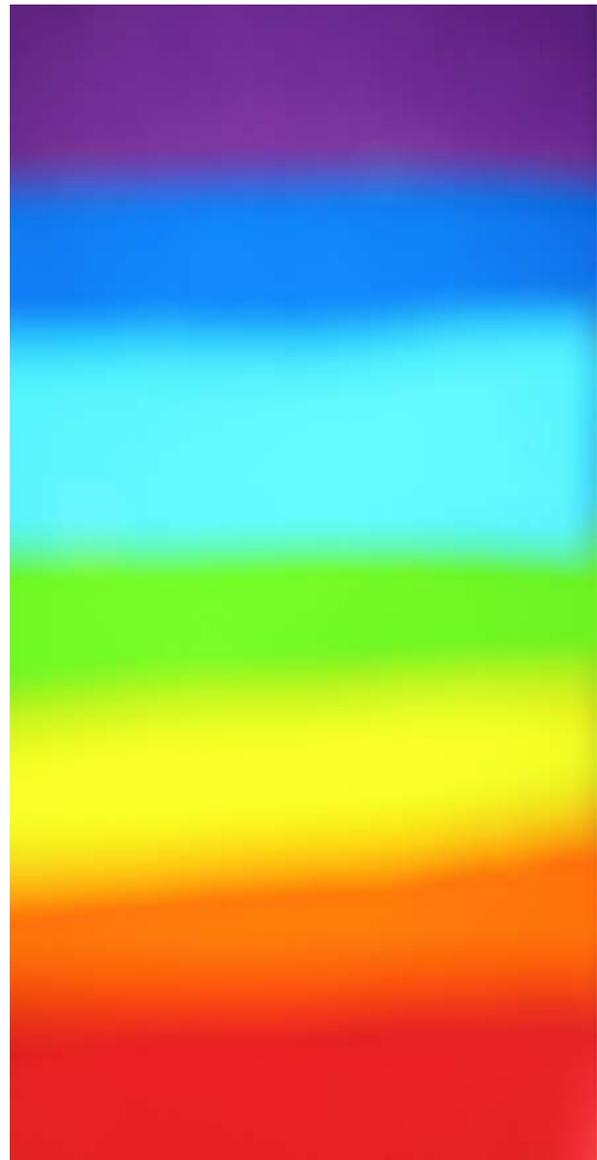
V.   PUBLIC SECTOR

How technology interfaces with, and has the power to impact ,historically marginalized communities should be a particularly heightened concern for public-sector organizations and policymakers. Governance bodies have a duty to ensure net societal benefits while protecting the public from harm. This means that, in the face of applying novel technology, balancing the potential for profound benefit while minimizing disparate and negative impacts should be the aim. More specifically, public policy should "make sure that people are not targeted, not harassed, and not murdered because of who they are, where they come from, who they love or how they pray."[2] The opportunity to model governance behaviors and practices at the highest level of accountability should also be considered.

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, to minimize the potential amplification of preexisting, discriminatory institutional structures. The Universal Declaration of Human Rights is a baseline, and its provisions should be prioritized above all else. Where other sectors and contexts fail to consider certain populations due to minority status, disenfranchised identity, cost effectiveness, 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 of individuals

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

The "Tragedy of the Commons" describes a phenomenon where 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, consider, for example, the commons of public privacy. 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—or rights—of the collective. These bounds should be informed by the values and priorities of the public, especially those most vulnerable, and apply to the principles and functions of public organizations.

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 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. This approach can have an outsized impact. One example is applying data science to long-term policy. For example, the policy on standard retirement age could be tied to median life expectancy. This creates policy that matures alongside society. Dismissing these approaches because they deviate from the norm is a missed opportunity. Governance bodies could be leveraging these capabilities to enshrine new policies that proactively iterate, while still allowing for intervention.

VI.   CONCLUSION

Leading organizations need to be intentional about their own behavior, and hold consideration for their impact that goes beyond government-mandated requirements. In doing so, companies have an opportunity to model responsible behavior, get out in front of competitors, and establish best practices and governance that can be codified and amplified by regulators. These are the companies that will set the bar for others to aspire to achieve. Will your company be setting the bar or playing catch up?

---

2   Sacha Baron Cohen, Recipient of ADL's International Leadership Award, Keynote Address at ADL's 2019 Never Is Now Summit on Anti-Semitism and Hate.

VII. AUTHORS

Steven Tiell

Nonresident Senior Fellow, GeoTech Center Atlantic Council

Steven Tiell is a nonresident senior fellow with the Atlantic Council's GeoTech Center. He is an expert in data ethics and responsible innovation working at Accenture, where he is the global lead for responsible innovation and data ethics. Since embarking on data ethics research in 2013, he has contributed to and published more than a dozen papers and has worked with dozens of organizations in financial services, government, high-tech, life sciences, media, public policy, public safety, and telecommunications sectors. He often speaks on topics such as governance, trust, data ethics, responsible innovation, surveillance, deepfakes, and industry trends.
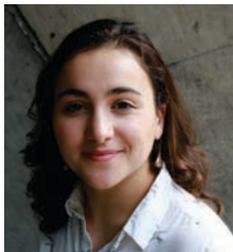
Lara Pesce Ares

Responsible Innovation Consultant Accenture

Lara Pesce Ares is an author of the Accenture Technology Vision. She develops thought leadership that covers technology futures and responsible business practices that often consider sociological implications. She is proud that her work creates impact through business-model innovations that position organizations to disrupt existing markets and enter new ones, influencing positive change at scale. She holds a BA in public policy from New York University, where she did a senior thesis on data-driven initiatives in city governments.

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