

Imagining Governance for Emerging Technologies

A new methodology from the National Academy of Medicine could inform social, ethical, and legal governance frameworks for a range of cutting-edge technologies.

Robyn is a 67-year-old Australian woman whose major depression was well managed with transcranial direct current stimulation (tDCS), a novel neurotechnology being used clinically in some countries, under the close supervision of her psychiatrist in Sydney. The treatment, which is approved by Australia's Therapeutic Goods Administration, was covered by her private insurance. However, when Robyn moved to the United States for work, she was dismayed to learn that the treatment is not approved here. Increasingly anxious about her worsening symptoms, she went online and ordered a "wellness aid" that looked something like the device she remembered from her doctor's office. One week and \$250 later, she began using the device at home, trying her best to replicate the way her psychiatrist placed it on her head and the settings she had used.

When his school day ends, Liam, a ninth grader, pops on the tDCS device that his parents got him for his birthday before he begins to play computer games. Although not regulated by the US Food and Drug Administration (FDA), the device he wears is marketed to adults as a way to improve attention and focus. Liam feels it helps him concentrate when he's gaming, as well as with his schoolwork and violin practice—so he uses it once or twice a day to try to maintain his edge at his competitive high school. His little brother, who is in third grade and struggles with attention-deficit/hyperactivity disorder, uses it occasionally as well.

These vignettes about Robyn and Liam, although hypothetical, highlight some of the questions raised by new tDCS technology. Is self-administered tDCS an adequate substitute for the mental health care that Robyn needs? What effects will chronic use have for Liam and his little brother? What other issues might arise from widespread and unregulated use of this emerging technology?

Transcranial direct current stimulation is one of many rapidly developing health-related technologies that transcend today's regulatory boundaries. As these technologies spread, they raise short-term questions like those embodied in these two vignettes, while also posing larger questions about how the technologies will affect society and how they should be governed. To begin to identify and address these questions for a range of new technologies, a new committee established by the National Academy of Medicine has developed a systematic methodology to inform a novel governance framework that considers not only the experiences of individuals like Robyn and Liam, but also anticipates larger social impacts.

This methodology is necessary because novel technologies developed to support and advance health and medicine that once stayed in the clinic are now making their way into workplaces, homes, entertainment, and beyond. Boundaries that once seemed clear—the difference between medical treatment and self-enhancement, for example, or the demarcation between therapy intended for adults and gaming devices intended for children—are blurring.

Although neural technologies such as tDCS may someday have the potential to transform mental health care, their ease of use, low cost, and relatively low physical risk have led to broad availability and use with little oversight.

By expanding access to promising technologies and driving their rapid evolution, such sector-crossing diffusion could be beneficial. But it can also broaden and magnify complex social, legal, and ethical issues. And as these technologies evolve simultaneously in multiple settings, it will be nearly impossible to anticipate or attend to their impacts on individuals, groups, or society at large.

How should such technologies be regulated and governed? It is increasingly clear that past governance structures and strategies are not up to the task. What these technologies require is a new governance approach that accounts for their interdisciplinary impacts and potential for both good and ill at both the individual and societal level.

To help lay the groundwork for a novel governance framework that will enable policymakers to better understand these technologies' cross-sectoral footprint and anticipate and address the social, legal, ethical, and governance issues they raise, our team worked under the auspices of the National Academy of Medicine's Committee on Emerging Science, Technology, and Innovation in health and medicine (CESTI) to develop an analytical approach to technology impacts and governance. The approach is grounded in detailed case studies—including the vignettes about Robyn and Liam—which have informed the development of a set of guiding principles (see sidebar).

Based on careful analysis of past governance, these case studies also contain a plausible vision of what might happen in the future. They illuminate ethical issues and help reveal governance tools and choices that could be crucial to delivering social benefits and reducing or avoiding harms. We believe that the approach taken by the committee will be widely applicable to considering the governance of emerging health technologies. Our methodology and process, as we describe here, may also be useful to a range of stakeholders involved in governance issues like these.

A boundary-crossing innovation ecosystem

Innovation today occurs in a vibrant ecosystem that features rapid technological development, adoption, and evolution. Added to that is the availability of capital, the push and pull of market incentives amid regulatory gaps, and convergence with other novel technologies and capacities. Although this environment creates opportunities for transformative, positive changes in health and medicine, it can also generate equally transformative harms that are far-reaching, difficult to anticipate, and even more difficult to address once manifest.

Returning to the example of tDCS: although the technology could yield an accessible treatment for depression or chronic pain, it could take a more fraught path outside

the medical context. Imagine, for instance, if the criminal justice system were to require prisoners to use a more advanced version of the device for behavior management.

Given that current regulatory regimes focus largely on single sectors—for example, the FDA regulates drugs and medical devices—there is a clear need to develop governance frameworks for emerging medical technologies that span multiple sectors. Furthermore, regulatory regimes for health and medicine generally attempt to forestall individual harms rather than understand and prevent potential society-wide harms. A framework for governance of developing technologies should intentionally drive toward societal benefit, instead of simply hoping it emerges from the market.

Considering how certain emerging technologies “muddle” efforts at governance, law and ethics scholars Gary Marchant and Wendell Wallach have written that these technologies “involve a complex mix of applications, risks, benefits, uncertainties, stakeholders, and public concerns. As a result, no single entity is capable of fully governing any of these

What these technologies require is a new governance approach that accounts for their interdisciplinary impacts and potential for both good and ill at both the individual and societal level.

multifaceted and rapidly developing fields and the innovative tools and techniques they produce.” Fostering the socially beneficial development of such technologies while also mitigating risks will require a governance ecosystem that cuts across sectors and disciplinary silos and solicits and addresses the concerns of many stakeholders.

And in this field of cross-sector technologies, governance itself is a complex ecosystem. The extent to which a technology's benefits are maximized and risks mitigated (and how benefits and risks are defined) often depends less on explicit ethical principles and values guiding the work itself, and more on the policies, norms, standards, and incentives of the particular sector that shapes a technology's development and deployment. A technology that develops within the sphere of clinical care will evolve in different ways than one that begins its trajectory in the gaming space.

Because every stage of the technology life cycle has a role in shaping what a technology will and will not do, sectoral governance takes many forms. These include formal governance through laws, regulations, executive orders, and court decisions, as well as more informal mechanisms such as norms and standards of conduct; guidance from scientific and medical academies or professional societies; and market forces such as consumer preferences.

Translating ethics to action

Building a novel governance framework that is up to the task of working in the fast-paced health and medicine innovation ecosystem requires a deeply interdisciplinary approach. The members of CESTI were intentionally selected to represent a broad array of sectors and interests within the technology ecosystem, including academic scientists and technologists, clinicians, bioethicists, and social scientists. Also included were technology investors and others from industry, who are frequently left out of such conversations due to conflicts of interest but offer critical insights regarding the values and incentives that drive and shape innovation.

The case study approach utilized by the committee is drawn from work one of us (Mathews) has done on developing analytic frameworks for emerging technologies. It is based largely on a set of features identified by the US Office of Technology Assessment in 1993 as common across many of the reports produced by the office, supplemented with the explicit consideration of ethical issues raised by the technology, the social goals of the research, and public engagement. The studies are designed so that reading several of them provides a bird's eye view of a technology in context. By revealing the myriad factors and interactions that shape the evolution and translation of these new technologies, the case studies can help us envision the elements of a cross-sectoral governance framework.

The committee selected three technologies to explore using the case study approach, each of which has a track record of governance as well as an anticipated trajectory of further evolution. The three technologies are tDCS, telemedicine, and regenerative medicine. Typical of the three is tDCS: although it has been studied in humans since at least the mid-1960s, improvements in microprocessors and batteries over the past ten years have coincided with a nearly tenfold increase in scholarly publications mentioning tDCS. The technology is expected to continue improving and evolving as domestic and international research programs reveal more about the brain and enabling technologies are further developed.

Likewise, over the past decade and particularly during the COVID-19 pandemic, interest in telemedicine has risen dramatically without clear guidelines for equitable use. This lack of guidance has persisted despite historical evidence that the expansion of telemedicine may be more likely to benefit those with privilege than those most in need.

Of the three technologies, regenerative medicine is unique in that it is a cellular technology focused on repairing or replacing damaged tissue and raises questions not only of financial access, but also biological access, since most of what researchers know in this field is based on the cells and genomes of people of European descent. Nonetheless, like the other two technologies, regenerative medicine has escaped the confines and regulatory structures of traditional academic research and clinical medicine, creating challenges that current governance systems are ill-equipped to handle.

Case studies

Each case study follows the same format. It starts with the technology and a sketch of the role it plays in people's lives today, followed by systematic investigation of the historical context, the status quo, a cross-sectoral footprint analysis, and exploration of ethical issues. The final section uses a scenario-based visioning exercise to build on the insights of the core case to explore how the technology might evolve and the role it could play in people's lives in the future. To provide a sense of the textured context that emerges from these case studies, we'll walk through the format using examples from tDCS.

All case studies begin with two short vignettes, like those of Robyn and Liam. Designed to make concrete a subset of the ethical issues raised by the case, the vignettes place the technology in realistic life stories so that the perspectives of individual stakeholders are an intrinsic part of each study.

The second section includes a brief introduction to the technology and an examination of its historical context, highlighting key scientific antecedents and ethics touchstones. For example, for any discussion of contemporary neurotechnologies such as tDCS, understanding the problematic history of psychosurgery is important. Thousands of vulnerable patients suffered significant harm, often in ways that reflected the sexism and racism of the broader society. The tDCS case study explains that in the 1960s, "reports that psychosurgery was being done not only on those with psychiatric disease (most of whom were women), but also on the young, including African American children," led Congress, the NIH, and the American Psychiatric Association to prepare a series of hearings, guidelines, and reports.

The case study's third section explores the status quo: identifying key questions surrounding the technology, outlining the current active areas of research, and highlighting available applications or products. Research results about the application of tDCS vary considerably across a wide range of study questions, including those related to learning and those related to physical performance, depression, and beyond, with some studies demonstrating modest effect and others demonstrating negligible or negative effects. Although there is some limited clinical use for treatment of major depression and chronic pain, only a few federal health authorities, such as in Australia and Singapore, have issued any level of device approval.

To systematically explore the footprint of a technology and how sectoral governance has shaped it, the fourth section of each case study is devoted to an extensive cross-sectoral analysis of how academia, health care, government, the private sector, and volunteers and consumers are engaged in its development. This analysis is broken down

into domains including science and technology; governance and enforcement; affordability and reimbursement; private companies; and social and ethical considerations. Although there is some overlap among the delineations of sectors and domains, this analysis has proven to be very revealing.

A broad array of governance mechanisms shapes the use of tDCS. These involve traditional regulators, such as the FDA, Consumer Product Safety Commission, and Federal Trade Commission, as well as professional societies and market forces. Particularly surprising to the committee was the way the current business models for tDCS devices affected the availability of venture capital for industry. Investors are somewhat reluctant to put money in commercial tDCS devices, in part because of difficulties with intellectual property in this space, but also because the device is likely to be a one-time purchase rather than a subscription involving regular payments—and there is already evidence that purchasers get bored quickly and abandon the contraption. Among the lessons learned from this observation is that commercial drivers can shape the development pathway of a technology even in the absence of

Pressure-testing future scenarios

The last section of each case study consists of a proactive and anticipatory visioning narrative designed to consider the future development of the technology, its applications, and its impacts on society. These narratives are the culmination of a visioning exercise in which committee members contributed their insights and expertise. To describe the environment or larger context within which an emerging technology is deployed, each narrative is constructed from the point of view of a relevant stakeholder—a family caregiver, health care worker, or businessperson, to name a few—giving a perspective of how the technology may evolve over the next 10 to 15 years. Ultimately, each visioning narrative should “pressure-test” the governance framework by investigating the questions: *Does the governance framework account for a range of plausible future scenarios; and if not, what needs to change?*

The visioning narrative for tDCS is written from the point of view of a fictional senior human resources executive at a software company in 2030, where use of tDCS in the workplace is as ubiquitous as the iPhone is today and “as unimaginable as the iPhone was in the 1990s,” writes the

By revealing the myriad factors and interactions that shape the evolution and translation of these new technologies, the case studies can help us envision the elements of a cross-sectoral governance framework.

efficacy data. We also learned that there has been significant interest from the military in using tDCS for performance enhancement.

Next, to better understand what is morally at stake with these emerging technologies, case studies include a broad list of ethical and societal implications highlighted by the technology, including whether particular applications raise unique equity concerns. For example, given the history of psychosurgery and its use in ways that reflected societal biases, special attention would be warranted regarding any use of tDCS or other neurotechnologies by law enforcement.

Building on these ethical implications, the next section, labeled “Beyond,” broadens the lens to consider the impact of other technologies that may facilitate or boost the technology in question. With tDCS, several brain-computer interface technologies are being developed to reduce tremors in Parkinson’s disease and enhance stroke recovery. Meanwhile, commercial developers are combining machine learning tools with measurements of users’ computer interactions to obtain digital proxies for their cognitive and emotional states. The combination of such capacities could catapult neurotechnologies into even more settings and applications, such as education and security, raising additional significant legal, ethical, and social questions.

executive. In the narrative, small brain-computer interfaces (BCIs) are used widely in medical care but have also become common as tools for better concentration, memory, and collaboration. The executive’s company has given the devices to more than 90% of employees. Noting that the devices significantly increased collaboration, the executive finds that they also raised questions about work-life balance. For example, “How does one separate personal thoughts from professional thoughts?” the executive wonders, before explaining that “we completely revamped the data governance around company owned BCIs after a foreign entity attempted to hack our chief financial officer’s device.”

Given the inherent uncertainties associated with the development trajectory of any emerging technology, the goal of this visioning exercise is not to accurately predict the future, but instead to create a plausible future scenario that surfaces latent ethical, legal, and regulatory tensions associated with the technology’s widespread use. In the visioning narrative above, those questions concern workplace applications, autonomy, privacy, and complex international security risks. Seeing these possibilities evolving in realistic scenarios encourages the exploration of a range of risks, opportunities, unintended consequences, and use contexts, enabling policymakers to set directions for governance.

Continued on page 46 →

Principles for the Governance of Emerging Technologies

Although there is a robust and growing literature on the ethics and governance of artificial intelligence, and another regarding gene editing technology, less has been said about the ethical principles that should guide the governance of emerging technologies in general.

To meet this need, the National Academy of Medicine's Committee on Emerging Science, Technology, and Innovation (CESTI) has taken a three-pronged approach that draws on cross-sectoral analyses of technological case studies, forward-looking visioning exercises, and a systematic ethical analysis. By examining these elements and identifying their points of convergence and divergence, CESTI has produced a set of ethical principles for the governance of emerging technologies that lays the groundwork for the translation of those ethical principles into policy decisions.

The CESTI principles are organized into three tiers: At the highest, most abstract level are ethics principles derived from the philosophical and bioethics literature. The middle tier includes more specific policy and governance goals, which are both procedural (speaking to how emerging technologies should be governed) and substantive (what that governance should achieve). The bottom tier, which remains a work in progress, consists of specific governance and policy tools that can be used to achieve those goals.

While the ethics principles in the top tier are more or less universal, the middle and bottom tiers were developed through conversation with the sectoral case analyses and the visioning exercises.

The topmost tier contains principles including justice, autonomy, fairness, collective good, and individual good. These principles derive from justice, autonomy, and beneficence, the three principles of biomedical ethics developed in the Belmont Report, a canonical text in the field of bioethics; we also reviewed consensus ethical principles and values that have been articulated for artificial intelligence and machine learning. CESTI has refined the content of these principles to be directly relevant to the governance of emerging technologies, acknowledging that a focus on discrete individual harms is insufficient given the reach and impact of these technologies.

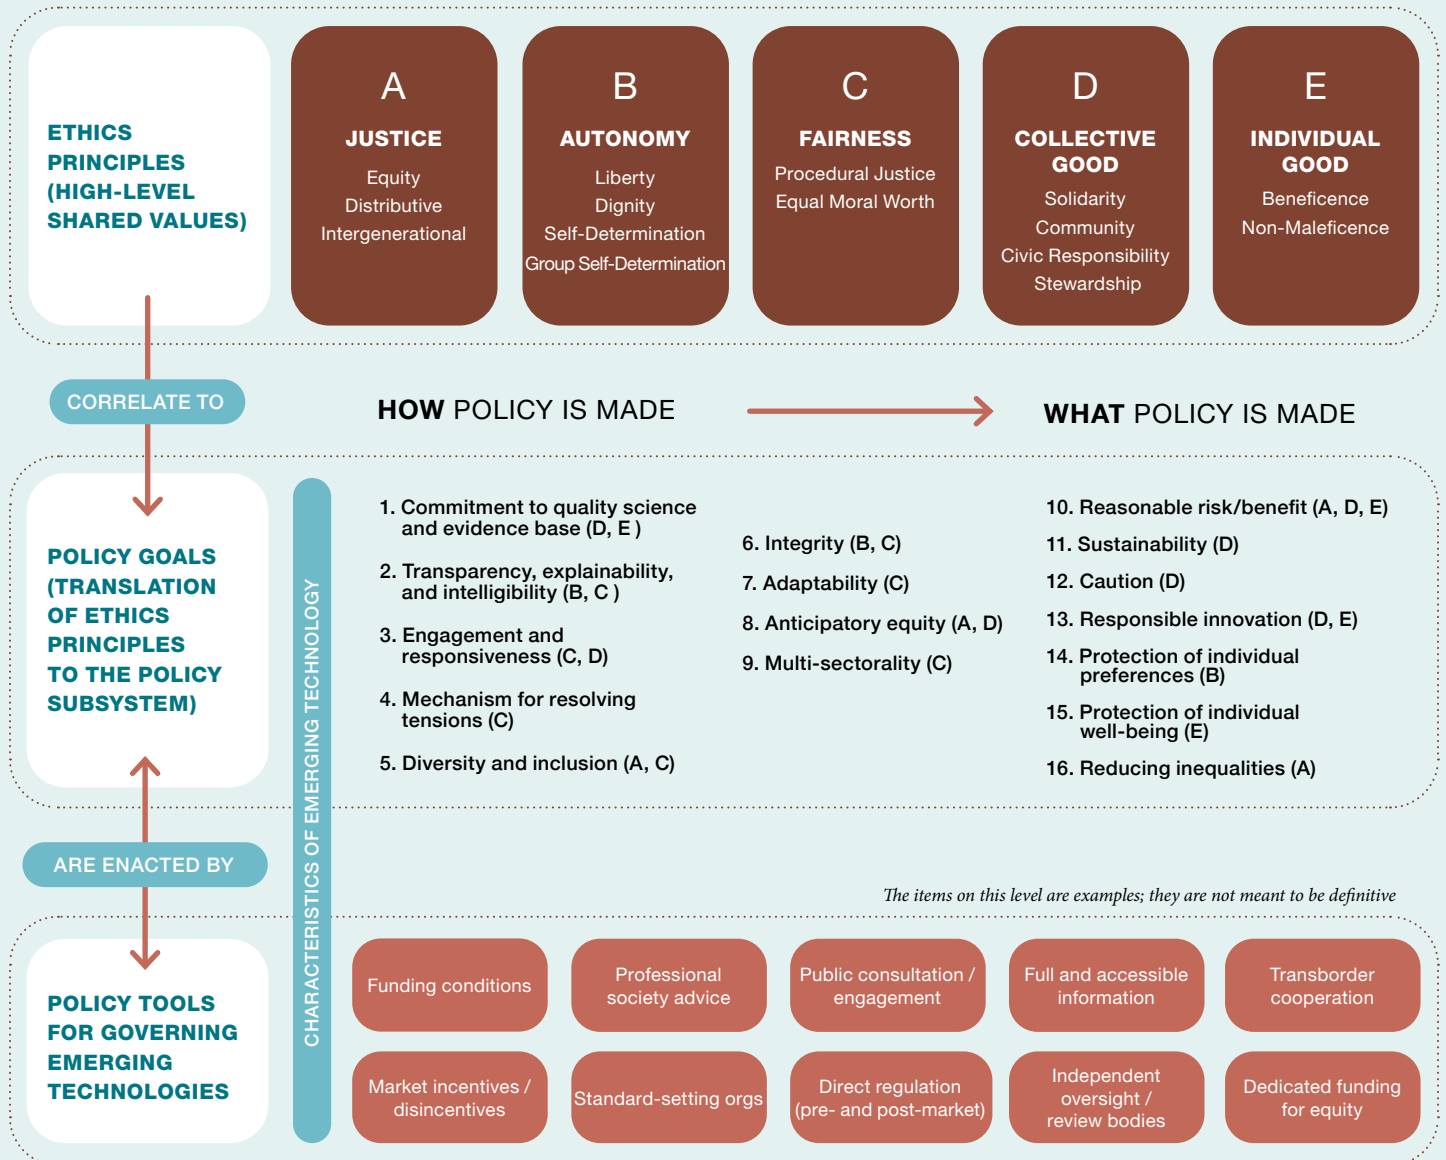
Justice, in the CESTI principles, refers to equity between groups faced with structural and systemic inequalities, a fair distribution of risks and benefits of technologies, and considerations about intergenerational justice, such as how decisions made now will affect future generations. Fairness refers to fair procedures for the creation of governance structures that are grounded in a view that all human beings are of equal moral worth, and may also reflect predictability and consistency, as well as transparency and accountability. Autonomy requires respect for individual

decisionmaking, liberty, and self-determination. Individual good requires that an emerging technology benefits individual users of that technology, and collective good requires the recognition that technologies have societal-level impacts (both benefits and harms) that are not captured by an exclusive focus on individuals.

After reflecting on these top-tier principles and learning lessons from the case studies, the committee articulated the policy and governance goals in the middle tier. This level identified 16 goals for governance that would be necessary for a system to uphold the top-level ethics principles. For instance, in the tDCS case, the committee recognized the challenge of governing technologies that have direct-to-consumer markets and do-it-yourself users that bypass most research oversight and federal regulation. To address this challenge, one goal, in the service of collective good and individual good, is ensuring that governance systems across sectors, including the private sector, are committed to quality science and an evidence base.

The development of the third tier of CESTI's translational approach, which considers governance tools and levers to achieve these goals, will be part of the work of the NAM consensus study that is taking up the CESTI mantle in 2022. With these principles to guide their work, the consensus study will have the groundwork necessary to ensure that the governance structures and frameworks they propose are firmly rooted in core ethical principles and values.

CESTI'S PRINCIPLES FOR THE GOVERNANCE OF EMERGING TECHNOLOGIES



The principles and the figure above were developed with CESTI's working group on Principles and Commitments: R. Alta Charo, I. Glenn Cohen, Jeffrey Kahn, Jennifer Kuzma, Gary Marchant, Prem Natarajan, Catherine Novelli, Kaushik Sunder Rajan, Jenny Reardon, and Keith Wailoo.

We began each visioning exercise by forming a working group of CESTI members with subject-matter expertise relevant to the technology in the case study. Outside subject-matter experts were also consulted. CESTI's visioning process builds on existing scenario-planning techniques to create a credible mid- and long-term perspective on an emerging technology. These perspectives may include elements such as rate of development, use cases, levels of public engagement and acceptance, regulatory requirements, and any larger societal forces or shifts that might shape future use. Relevant societal shifts might include a significant increase in remote working, erosion of public trust in scientific institutions, greater societal emphasis on work-life integration, and the development of health care tourism.

We interviewed the working group members and outside experts to create a shared understanding of the current “baseline” performance of an emerging technology. Through these discussions, we identified key assumptions around the technology's maturation, mapped out a potential future application with clear implications for health, and identified any key societal shifts that might characterize the general environment within which the technology was adopted.

Following this level-setting and development of a rough trajectory, we established the overarching theme of the scenario. With tDCS, we focused on a miniaturized brain-computer interface that significantly accelerates and reduces friction for remote work and workplace communication.

The intention behind these scenarios is to improve the rigor of decisionmaking and strategic planning with respect to an intended governance framework by providing a rich, detailed accounting of how a technology and its social context may evolve. To do this, visioning narratives must be plausible, anchored to the perspective of a major stakeholder who is likely to influence or be influenced by the issues identified, and internally consistent with respect to time frame, stakeholder perspective, and use case. In the case of the brain-computer interface, the narrative included workplace discrimination, exacerbation of income inequality, addictive behavior around its use, and other issues.

Opportunities for a more engaged system of governance

The methodology of using core cases provides the understanding and scaffolding necessary for decision-makers to identify the critical elements of a novel governance framework. The power of the visioning narrative lies in its ability to describe the larger context in which an emerging technology is governed. We believe that these case studies, and the process we used, will prove helpful not only for policymakers and regulators, but also for industry and academic explorations of innovation governance.

Further, there are important insights to be gained by looking at multiple case studies to inform the development of a cross-sectoral governance framework intended to be shaped and guided by a set of overarching principles (see sidebar). Gone are the days when a single regulator or a siloed governing framework can oversee medical and health technologies.

Looking to the future, we see many opportunities for a more engaged and integrated system of governance that strives to orient technologies toward social benefits. Our approach makes an argument for and gives insights into how policy changes could shift social outcomes. For example, knowledge of and visioning about unequal access to regenerative medicine demanded the inclusion of justice as a guiding ethical principle, reflected in the policy goal of “anticipatory equity.” In addition, we learned from the regenerative medicine case that if a governance structure is not designed to consider equity questions early in a technology's life cycle, it may be considerably more difficult to address them later—as can be seen today, with the development of telemedicine outpacing equitable access to broadband internet. A governance structure that could anticipate this potential challenge through fair and inclusive procedures and attention to structural injustice might prevent similar equity concerns from arising in the future.

In this complicated clinical and commercial innovation ecosystem, we think that the CESTI approach will help build a foundation that will eventually enable policymakers to understand the complexities of regulating emerging technologies and make policy choices to safeguard both individuals and society. We're making the building blocks of our approach available to practitioners from many fields in the hope that they will be useful in related efforts.

Debra J. H. Mathews is the assistant director for science programs for the Johns Hopkins Berman Institute of Bioethics and an associate professor in the Department of Genetic Medicine at Johns Hopkins University School of Medicine.

Rachel Fabi is an assistant professor in the Center for Bioethics and Humanities at the State University of New York Upstate Medical University. Anaeze C. Offodile II is an assistant professor in the Department of Reconstructive Surgery and the enterprise executive director for clinical transformation at MD Anderson Cancer Center. The authors are writing under the auspices of the National Academy of Medicine's Committee on Emerging Science, Technology and Innovation in health and medicine, which is cochaired by Atul Butte, R. Alta Charo, and George Q. Daley. The authors wish to acknowledge the contributions of the full committee, a roster of which can be found here: <https://nam.edu/programs/committee-on-emerging-science-technology-and-innovation-in-health-and-medicine/>.