

Mobilizing the private sector for responsible innovation in neurotechnology

To the Editor — Emerging neurotechnologies raise important governance questions related to, for example, dual use, brain data privacy, and manipulation of personal autonomy. Although many public sector research initiatives have implemented measures to address these issues, similar systematic measures in the private sector have yet to emerge. This gap is critical, as neurotech innovation today is largely driven by a set of companies that are subject to growing public scrutiny^{1–5}. Here we detail lessons, emerging practices and open questions for responsible innovation in the private sector that are the result of three years of policy deliberations that began with a 2018 conference in Shanghai convened by the Organization for Economic Co-operation and Development (OECD) and led to the release of the “OECD Recommendation on Responsible Innovation in Neurotechnology” last year⁶. The principles therein cover opportunities and challenges for better innovation practices in company settings—including the use of ethics advisory boards, company-level principles, and ethics-by-design approaches—with broad relevance beyond neurotech to digital medicine and corporate R&D activities in today’s era of ‘tech-lash’. We argue that it is time for a radical shift in the conversation about governance of emerging neurotech: effective governance must focus on the private sector as a central actor early on—before trajectories are locked in and scaling takes off—and requires a new set of policy perspectives and collaborative tools to do so. These tools must complement existing efforts in public-sector research ethics, post hoc product regulation and corporate social responsibility. They must also reflect the growing recognition that we cannot rely on industry self-regulation alone to steer innovation activity in socially desirable directions.

Missing the mark with ethics and responsible innovation initiatives?

Emerging neurotechnologies, such as brain–computer interfaces (BCIs) or digital phenotyping for mental health monitoring, hold considerable promise for health and well-being, but also raise important ethical, social, and governance questions (Box 1). These questions include concerns about brain data privacy, runaway human enhancement, individual autonomy, vulnerability to

political or economic manipulation, direct-to-consumer (DTC) marketing of devices that have variable, if any, effectiveness, dual use, do-it-yourself (DIY) neurotech and neurohacking, and new forms of inequality^{7–11}. Although public-sector research has been quick to implement targeted programs to tackle these concerns—for example, the ‘Ethics and Society’ strand of the Human Brain Project^{12,13}—the private sector has thus far paid relatively scarce systematic attention to them³.

This gap between public and private sector efforts to foster responsible innovation practices is critical insofar as, across OECD countries, over 70% of all R&D is performed by the private sector. Moreover, the implications of many recent innovations have become fully visible to society only when scaled up by companies. This has placed some of the most successful technology firms increasingly into the crosshairs of regulators and a public tech-lash. Facebook, for instance, has been subject to a barrage of inquiries about the free speech and content moderation, data protection, or the effects of surveillance capitalism and echo chambers on democracy. Clearview AI, a facial recognition software company, promises greater public safety through scalable, app-based facial-recognition techniques, but has been criticized for enabling intrusive and potentially authoritarian uses. Digital platforms that have begun to transform entire service sectors, such as Uber or Airbnb, have also raised concerns about new inequalities pertaining to undermining labor laws or driving real estate speculation, respectively. From a policy perspective, these big tech examples beg the question of whether ‘responsible innovation’ efforts focused on public-sector research simply miss the mark. The same holds true for emerging neurotechnologies.

Traditional technology governance is increasingly insufficient

For emerging neurotech, traditional means of governance—including institutionalized research ethics, post hoc regulation and market mechanisms—are ill-equipped to capture the ways in which these technologies could reshape our societies, especially in terms of long-term consequences. The potential uses of non-invasive BCIs in the workplace, for example, are raising new

controversies about labor protection and employee surveillance¹⁴. Likewise, there is a debate as to whether research into certain types of BCI should be banned because of dual-use applications (for example, covert manipulation of personality), thus foregoing potential civil-use benefits (for example, the restoration of sensorimotor functions after spinal cord injury)^{10,15}. Challenges may even rise to the judicial and constitutional level: the landmark federal case *US v. Semrau* for the first time considered, though ultimately dismissed, brain scans as a source for lie detection¹⁶. The Chilean senate is considering a constitutional amendment that, if approved, would be the first to legally codify ‘neurorights’ to protect the mental integrity and privacy of its citizens¹⁷. This reflects wider debates about the need for new human rights in the age of rapidly evolving neuroscience and neurotech¹⁸.

Neurotech, like many other innovation domains, is also subject to a patchwork of national and regional regulations that create considerable uncertainty. National attempts to govern emerging technology are frequently seen as ineffective or even detrimental to innovative economies, prompting concerns that companies and technologies may simply move across borders. Developing new international treaties, however, is notoriously difficult, and intergovernmental organizations often rely on soft law, such as the OECD recommendation^{6,19}. Within single jurisdictions, too, there is ample regulatory complexity, as neurotech straddles sectors, applications and regulatory domains. BCIs, digital phenotyping apps and psychopharmacology will be subject to different regulatory regimes straddling health, safety, trade or drug regulation, each of which are governed by different governmental agencies and jurisdiction.

Recognizing these challenges, policymakers have increasingly tried to engage upstream governance approaches—that is, early interventions during the research process—to complement traditional post hoc regulation. In public-sector research, approaches such as anticipatory governance²⁰ and responsible research and innovation (RRI)²¹ have gained increasing credibility. Instruments of Ethical, Legal and Social Implications (ELSI) piloted by the Human Genome Project²², such as

Box 1 | Responsibility issues in neurotech companies

- **Brain data privacy.** Companies developing products, services and treatments relying on neurodata, whether from behavioral apps or BCIs, face mounting questions regarding data ownership, security, privacy and consent. This debate echoes similar concerns raised in genomics or digital platforms, but involves potentially more sensitive data around individual motivations that could be used in court or for unintended purposes.
- **Human enhancement.** Many companies are developing products that walk a fine line between rehabilitation (for example, for war veterans' injuries) and human enhancement (for example, conditioning). Echoing long-standing debates about psychopharmacology, this tension raises questions about different national standards for human enhancement, especially when combined with BCI technology and brain stimulation. Similar concerns have been raised in the context of novel gene editing techniques.
- **Autonomy and vulnerability to manipulation.** Many emerging app-based neurotechnologies exploit behavioral patterns to predict or respond to cognitive states. The data collected by these apps may potentially be used to trigger mental health interventions, but can make individuals also vulnerable to surveillance, policing, and economic or political manipulation. Examples include targeted advertisements or behavioral nudges. Some emerging technologies have raised concerns about workplace surveillance—for example, measuring attention in the context of digital work environments.
- **Direct-to-consumer marketing.** Marketing as a DTC technology for well-being and entertainment is the pathway of choice for many startups trying to avoid the high regulatory scrutiny associated with medical devices. These boundaries can be blurred easily, however, and many companies would prefer clear guidelines over amorphous expectations.
- **Dual use.** Many products, especially those concerned with neuroenhancement or suppression of cognition, straddle medical, consumer and military applications. For companies, this situation raises questions about export markets and potential security concerns.
- **Do-it-yourself neurotech and neurohacking.** There is a growing DIY neurotech scene similar to that in other biotechnology domains, including gene editing. While sometimes welcomed as a sign for a vivid entrepreneurial ecosystem, neurohackers may fly under the regulatory radar and develop potential products or applications that might otherwise attract high scrutiny.
- **New forms of inequality.** All of the aforementioned challenges may introduce new forms of inequality. Unequal access to medical treatments, populations vulnerable to exploitation (for example, children or those unable to understand terms of service), a 'race to the bottom' for enhancement, or uneven regulation across countries are examples. Many companies are gauging the global landscapes for different business models.

focus group research, citizen juries, ethical review through institutional review boards, or stage-gate processes, have been taken up by the Human Brain Project and BRAIN Initiative, among others²³.

Although ELSI and RRI frameworks have successfully penetrated wide parts of public research on neurotech, similar systematic frameworks in the private sector are lacking. Companies tend to sit in a blind spot between early-stage research ethics and post hoc regulatory responses that primarily focus on safety and efficacy, monopoly power or liability. Approaches found in RRI or ELSI programs are neither mandatory nor easily applicable in corporate settings. For one, the broader social consequences of technological change—including

new forms of inequality, vulnerability or risk—are hard to capture as part of company metrics, incentive structures and shareholder value logics. For another, the field is largely driven by startup dynamics, which does not afford extended time for deliberation or dedicated organizational resources. The entrepreneurial mindset to move fast, break things, scale up, and worry about consequences later²⁴ is at odds with traditional governance mechanisms such as ethics board reviews and public consultations during product development. This need for speed and scale can lead to unintended consequences as well as overpromising. For example, Lumosity, a company providing a brain-training app, was fined \$2 million by the US Federal

Trade Commission in 2016 for deceptive claims about enhanced concentration and decreased cognitive impairment in patients with Alzheimer disease using its products. In the absence of good strategies, governments are increasingly embracing experimental strategies to tackle governance challenges or test applications. The US Food and Drug Administration is testing experimental precertification programs partly to get a grip on, among other things, emerging mobile applications for mental health that are increasingly marketed DTC. The city of Reno has embarked on a local experiment to offer app-based mental health services for its residents through the company Talkspace to help alleviate the devastating mental health effects of the COVID-19 pandemic. This took place despite recent controversy around privacy issues for apps, reflecting a common 'hands-off' approach by local jurisdictions towards responsible innovation.

What companies should do

The current lack of systematic responsibility frameworks does not mean that embedded upstream governance options for emerging technologies cannot be implemented in the private sector. Our three-year dialogue processes revealed that a range of neurotech companies are actively seeking guidance and developing their own toolkits to bridge structural constraints and the apparent need for greater public oversight. What is more, many leading neurotech companies have a strong interest in publicly demonstrating responsibility and integrity, recognizing that the entire nascent sector can be harmed by single irresponsible actors in the field. Below, we list several emerging practices and principles that can help ensure better governance of neurotechnology innovation in corporate settings.

Enable responsibility review and diverse perspectives as part of the R&D process.

One example of a company that appointed an ELSI Advisory Board early in its history is Mindstrong, a company that develops apps to predict mental illness relapses through patient smartphone interactions. This board brought together engineers, ethicists, social scientists and people living with mental health issues to actively shape development of the technology. It was instrumental, for example, in the decision to switch from collecting text or global positioning system (GPS) data, which users considered intrusive of their privacy, to content-free and less readily identifiable signals from the smartphone, such as keyboard interactions patterns²⁵. This diversity-oriented advisory board strategy is broadly consistent with the

recent surge in corporate hires from the humanities and social sciences to inject critical and socially inclusive perspectives into innovation processes. Getting these structures right and sustaining them in a corporate environment is not trivial, however. Google famously had to dissolve its AI Ethics Council just one week after its much-anticipated launch, following considerable internal and external backlash about its composition.

Develop robust responsibility principles as part of a startup's mission. One of us (D.B.) has developed a code of responsibility for his neurotech startup Aifred, which applies deep-learning algorithms to enhance individualized psychiatric treatment. In this 'meticulous transparency' framework, all machine-learning projects must be reviewed by the clinical and machine-learning team with respect to their intended outcome, the target population, the representativeness of the available data, interpretability metrics, and monitoring for adverse effects of the model²⁶. The framework helped resolve concrete design dilemmas like the use of binary predictive algorithm outputs, such as 'being' or 'not being' at risk of suicide—which the company decided should best be designed as a warning system available only to clinicians and only with probabilistic, rather than binary, outputs. This, in turn, affected the way the machine-learning analyses were conceptualized—an example of responsibility-driven design. The focus on responsibility as part of concrete, embedded code differs from the rather high-level, non-committal ethics guidelines for artificial intelligence and other technologies released by the dozens by many corporate giants.

Embrace collectively legitimated ethics-by-design approaches. Standard-setting bodies like the Institute of Electrical and Electronics Engineers (IEEE) are increasingly targeting the engineering phase of product development to address social values and standardize certain critical features from the beginning, including in the fields of neurotech and artificial intelligence²⁷. Upstream ethics-by-design approaches aim to hardwire values into downstream developments. Given their consequences, however, these choices need to be opened up to collective deliberation and be subject to some form of political legitimation. Through bodies like the IEEE, public- and private-sector actors can work together to collectively define product standards and codify responsible design choices that embody shared commitments around values such as privacy and transparency.

Mobilize tech transfer as a critical juncture for social impact. Many universities are adjusting their tech transfer rules to better reflect social priorities. They are emphasizing, for example, inclusiveness in benefit-sharing and requirements to institutionalize certain values and accountability structures. Historically, the incentive structures for technology transfer offices have tended to maximize revenue, the number of startups, or scope of corporate sponsorship, with little attention to ethical and social deliberation elements. With an impressive list of signatories, the "Nine Points to Consider" code of good practice in university technology transfer provides a model for how to leverage technology transfer for more responsible innovation practices, including in neurotech²⁵.

Pressure investors to select for responsible technology development approaches. Shareholders are increasingly stepping up to inject responsibility considerations in company strategy. In 2018, two major Wall Street investors pressured Apple to take steps to fight addictions to iPhone use in children, which led to the development of an app called Screen Time. In neurotech, some companies are actively seeking out investors who match their values. Yet, as the Shanghai OECD conference revealed, the number of venture capital investors foregrounding responsible innovation concerns is limited, despite considerable interest among startups to work with specialized investors who know and acknowledge the ethical and social challenges of their technologies. This opens up an opportunity for a new subset of investment instruments or venture capital niches dedicated to responsible innovation practices, similar to the recent surge in sustainable investment and 'green bond' portfolios that target environmental or climate-related projects²⁸. Such developments could be further supported by new standards or certifications on responsible investment in tech startups.

Rethink corporate social responsibility approaches. Traditional corporate social responsibility (CSR) typically addresses the protection of workers, local communities and the environment through self-governance tools. However, CSR has largely ignored innovation as a key arena for social impact and responsible business conduct—as evident in today's controversies surrounding 'big tech'. In most neurotech companies, CSR approaches do not help solve the aforementioned ethical, social and governance dilemmas. Likewise, engineering ethics frameworks tend to remain outside the purview of CSR²⁹. Targeting the next

generation of innovators, a growing number of universities are offering resources for students, entrepreneurs and startups to consider responsibility and risks to sustainability as part of business model development, including Arizona State University's Risk Innovation Nexus and the Technical University of Munich's Master of Arts program "Responsibility in Science, Engineering, and Technology," in which one of us (S.P.) is involved. The incorporation of responsible innovation into engineering education and nascent business models can create an added value—for example, by gauging long-term societal implications or engaging early with potential future concerns or regulations-in-the-making. In the long run, the disconnect between CSR and corporate R&D raises serious questions about the adequacy of traditional CSR approaches for an era in which business models tend to focus on innovation and disruption.

Finding the right balance

The past 10 years have brought into sharp relief not only seemingly unregulated spaces in which innovative companies can rapidly grow from small startups to powerful global forces, but also the difficulties in exerting regulatory scrutiny in real-time through traditional governance approaches. The burgeoning field of neurotech is no exception. Yet, cognizant of the fallout recently observed in the controversies surrounding 'big tech', many neurotech companies are actively looking for guidance on how to increase the social robustness and sustainability of their emerging products and services in this field.

There is, of course, reason to be skeptical that companies alone will ensure socially responsible technology trajectories. Industry self-regulation has regularly failed to deliver the promised results and has instead stoked critiques of tokenization and greenwashing. A similar effect can arguably be observed in the current wave of 'ethics washing' (the practice of implementing superficial ethics mechanisms or principles in response to public pressure while purposefully side-stepping more fundamental issues, as most recently observed in the controversies surrounding Facebook's Oversight Board). Many large tech companies have accepted the cost of asking for forgiveness—such as paying out fines or legal settlements—as the cost of doing business. Thus, the promise of more responsible innovation by way of industry self-governance can only supplement, rather than supplant, public oversight.

However, there is ample evidence that government regulation alone will not suffice, as traditional policy instruments are increasingly at a disadvantage in today's

innovation landscape. This is why dedicated forums and spaces that can mobilize alliances of companies, policymakers, academics and citizens are needed to raise the bar for responsible innovation and co-develop new mechanisms for self-governance on the business side alongside new government regulation, including the ones mentioned above. International organizations such as the OECD or IEEE are uniquely positioned to speak to differences in national regulations and are already playing key roles in fostering the necessary dialogues^{3,6,27}. Universities, too, can mobilize their educational and entrepreneurial ecosystems to heighten sensitivity to responsibility concerns and foster policy dialog when companies are in the startup stage. What is more, experimental 'living lab' and 'sandbox' approaches could be used to co-develop new regulations and foster public debate about novel technologies, not just to create pro-business innovation environments through lower regulatory standards, as is currently the case in many such settings³⁰. Neurotech companies, with obvious social and ethical challenges on the horizon, have a chance to set an example for the entire tech industry. □

Sebastian M. Pfotenhauer¹✉, Nina Frahm¹, David Winickoff², David Benrimoh^{3,4}, Judy Illes⁵ and Gary Marchant⁶

¹Munich Center for Technology in Society (MCTS) and TUM School of Management, Technical University of Munich, Munich, Germany. ²Working Party for Bio-, Nano- and Converging Technologies, Organisation For Economic Co-operation and Development (OECD), Paris, France. ³Department of Psychiatry, McGill University, Québec, Montréal, Canada. ⁴Aifred Health, Montreal, Quebec, Canada. ⁵Neuroethics Canada, Department of Medicine,

University of British Columbia, Vancouver, British Columbia, Canada. ⁶Sandra Day O'Connor College of Law, Global Institute for Sustainability and Innovation, Arizona State University, Phoenix, AZ, USA.

✉e-mail: sebastian.pfotenhauer@tum.de

Published online: 7 June 2021
<https://doi.org/10.1038/s41587-021-00947-y>

References

- Eaton, M. L. & Illes, J. *Nat. Biotechnol.* **25**, 393–397 (2007).
- Jarchum, I. *Nat. Biotechnol.* **37**, 993–996 (2019).
- Garden, H., Winickoff, D., Frahm, N. & Pfotenhauer, S. M. Responsible innovation in neurotechnology enterprises. OECD Science, Technology and Industry Working Papers, No. 2019/05 (OECD, 2019); <https://doi.org/10.1787/9685e4fd-en>
- Wexler, A. A skeptic's take on Neuralink and other consumer neurotech. *Stat* (7 April 2021).
- Wexler, A. & Reiner, P. B. *Science* **363**, 234–235 (2019).
- OECD. *OECD Recommendation on Responsible Innovation in Neurotechnology* <https://www.oecd.org/science/recommendation-on-responsible-innovation-in-neurotechnology.htm> (2019).
- Wexler, A. *Front. Hum. Neurosci.* **11**, 224 (2017).
- Schwartz, A. *Science* **350**, 11 (2015).
- Ienca, M., Haselager, P. & Emanuel, E. J. *Nat. Biotechnol.* **36**, 805–810 (2018).
- Ienca, M., Jotterand, F. & Elger, B. S. *Neuron* **97**, 269–274 (2018).
- Nuffield Council on Bioethics. *Novel Neurotechnologies: Intervening in the Brain*. <https://www.nuffieldbioethics.org/publications/neurotechnology> (2013).
- Salles, A. et al. *Neuron* **101**, 380–384 (2019).
- Global Neuroethics Summit Delegates. et al. *Neuron* **100**, 19–36 (2018).
- De Stefano, V. Neuro-surveillance and the right to be human at work. *On Labor* <https://onlabor.org/neuro-surveillance-and-the-right-to-be-humans-at-work/> (15 February 2020).
- Clausen, J. et al. *Science* **356**, 1338–1339 (2017).
- Jones, O. D., Marois, R., Farah, M. J. & Greeley, H. T. *J. Neurosci.* **33**, 17624–17630 (2013).
- Asher-Schapiro, A. Out of my mind: advances in brain tech spur calls for 'neuro-rights'. *Reuters* (29 March 2021).
- Ienca, M. & Andorno, R. *Life Sci. Soc. Policy* **13**, 5 (2017).
- Marchant, G. & Allenby, B. *Bull. At. Sci.* **73**, 108–114 (2017).
- Barben, D., Fisher, E., Selin, C. & Guston, D. in *The Handbook of Science and Technology Studies* (eds Hackett, E. J. et al.) 979–1000 (MIT Press, 2008).
- Stilgoe, J., Owen, R. & Macnaghten, P. *Res. Policy* **42**, 1568–1580 (2013).
- NIH. ELSI planning and evaluation history. <https://www.genome.gov/10001754/elsi-planning-and-evaluation-history> (2012).

- Frahm, N. M. International governance of neuroscience and neurotechnology: whom to trust with the assessment of future pathways? *The Neuroethics Blog* <http://www.theneuroethicsblog.com/2017/11/international-governance-of.html> (2017).
- Pfotenhauer, S.M., Laurent, B., Papageorgiou, K. & Stilgoe, J. *Soc. Stud. Sci.* (in the press).
- AUTM. *In the Public Interest: Nine Points to Consider in Licensing University Technology* https://www.autm.net/AUTMMain/media/Advocacy/Documents/Points_to_Consider.pdf (2007).
- Benrimoh, D., Israel, S., Perlman, K., Fratila, R. & Krause, M. in *Recent Trends and Future Technology in Applied Intelligence* (eds Mouhoub, M. et al.) 869–880 (Springer International, 2018); https://doi.org/10.1007/978-3-319-92058-0_83
- IEEE. *Ethically Aligned Design, Version 1, Translations and Reports* <https://standards.ieee.org/industry-connections/ec/ead-v1.html> (2016).
- Kurtz, E. in *The Oxford Handbook of Corporate Social Responsibility* (eds Crane, A. et al.) 249–280 (Oxford Univ. Press, 2008); <https://doi.org/10.1093/oxfordhb/9780199211593.003.0002>
- Smith, N. M., Zhu, Q., Smith, J. M. & Mitcham, C. *Sci. Eng. Ethics* **27**, 28 (2021).
- Engels, F., Wentland, A. & Pfotenhauer, S. M. *Res. Policy* **48**, 103826 (2019).

Acknowledgements

S.P. acknowledges support by the European Union's Horizon 2020 research and innovation program under grant agreement 788359 ('SCALINGS: Scaling up co-creation? Avenues and limits for including society in science and innovation'). J.I. is supported by the Canada Research Chairs Program as the Canada Research Chair in Neuroethics and co-leads the Canadian Brain Research Strategy (CIHR Grant #171583/03027 IC-127354).

Competing interests

S.P. serves as academic program director for the master's program 'Responsibility in Science, Engineering, and Technology' (REST) at Technical University of Munich. D.W. is a senior policy analyst at the Organisation for Economic Co-operation and Development (OECD), where he heads the Working Party for Biotechnology, Nanotechnology and Converging Technologies (BNCT) responsible for the Recommendation on Responsible Innovation in Neurotechnology. D.B. is a founder, shareholder and employee of Aifred Health, a digital mental health company.

Additional information

Peer review information *Nature Biotechnology* thanks the anonymous reviewers for their contribution to the peer review of this work.



The need for global access to biomedical innovations during pandemics

To the Editor — The COVID-19 pandemic has fundamentally reshaped the way in which research and development is conducted and, in the case of vaccines and therapeutics, compressed a decade-long product development process to less than one year¹. However, the new technologies produced by this accelerated innovation—such as mRNA vaccines, viral vector vaccines, monoclonal antibodies and state-of-the-art diagnostics—have not been equitably distributed worldwide. Intense

early competition among high-income and upper-middle-income countries over initially scarce supplies has prevented low- and middle-income countries (LMICs) from accessing these innovations². For vaccines, this problem will likely continue if annual vaccinations are required and/or SARS-CoV-2 variants warrant modification and subsequent boosters—both requiring additional supplies. The lack of delivery systems to distribute these innovations in many LMICs further compounds this

situation^{3,4}. At best, most populations in LMICs will face a protracted delay of months to years in receiving these innovations at scale; at worst, some populations may never receive them^{5,6}.

To effectively respond to pathogens with pandemic potential, the product development ecosystem needs to both develop biomedical products in an even shorter timeframe and distribute these products globally in an equally rapid manner. Building on [proposed strategies](#)