

Regulating Automated Healthcare and Research Technologies

First Do No Harm (to the Commons)

Roger Brownsword

27.1 INTRODUCTION

New technologies, techniques, and tests in healthcare, offering better prevention, or better diagnosis and treatment, are not manna from heaven. Typically, they are the products of extensive research and development, increasingly enabled by high levels of automation and reliant on large datasets. However, while some will push for a permissive regulatory environment that is facilitative of beneficial innovation, others will push back against research that gives rise to concerns about the safety and reliability of particular technologies as well as their compatibility with respect for fundamental values. Yet, how are the interests in pushing forward with research into potentially beneficial health technologies to be reconciled with the heterogeneous interests of the concerned who seek to push back against them?

A stock answer to this question is that regulators, neither over-regulating nor under-regulating, should seek an accommodation or a balance of interests that is broadly ‘acceptable’. If the issue is about risks to human health and safety, then regulators – having assessed the risk – should adopt a management strategy that confines risk to an acceptable level; and, if there is a tension between, say, the interest of researchers in accessing health data and the interest of patients in both their privacy and the fair processing of their personal data, then regulators should accommodate these interests in a way that is reasonable – or, at any rate, not manifestly unreasonable.

The central purpose of this chapter is not to argue that this balancing model is always wrong or inappropriate, but to suggest that it needs to be located within a bigger picture of lexically ordered regulatory responsibilities.¹ In that bigger picture, the paramount responsibility of regulators is to act in ways that protect and maintain the conditions that are fundamental to human social existence (the commons). After that, a secondary responsibility is to protect and respect the values that constitute a group as the particular kind of community that it is. Only after these responsibilities have been discharged do we get to a third set of responsibilities that demand that regulators seek out reasonable and acceptable balances of conflicting legitimate interests. Accordingly, before regulators make provision for a – typically permissive – framework that they judge to strike an acceptable balance of interests in relation to some particular technology, technique or test, they should check that its development, exploitation, availability and application crosses none of the community’s red lines and, above all, that it poses no threat to the commons.

¹ See, further, R. Brownsword, *Law, Technology and Society: Re-imagining the Regulatory Environment* (Abingdon: Routledge, 2019), Ch. 4.

The chapter is in three principal parts. First, in Section 27.2, we start with two recent reports by the Nuffield Council on Bioethics – one a report on the use of Non-Invasive Prenatal Testing (NIPT),² and the other on genome-editing and human reproduction.³ At first blush, the reports employ a similar approach, identifying a range of legitimate – but conflicting – interests and then taking a relatively conservative position. However, while the NIPT report exemplifies a standard balancing approach, the genome-editing report implicates a bigger picture of regulatory responsibilities. Second, in Section 27.3, I sketch my own take on that bigger picture. Third, in Section 27.4, I speak to the way in which the bigger picture might bear on our thinking about the regulation of automated healthcare and research technologies. In particular, in this part of the chapter, the focus is on those technologies that power smart machines and devices, technologies that are hungry for human data but then, in their operation, often put humans out of the loop.

27.2 NIPT, GENOME-EDITING AND THE BALANCING OF INTERESTS

In its report on the ethics of NIPT, the Nuffield Council on Bioethics identifies a range of legitimate interests that call for regulatory accommodation. On the one side, there is the interest of pregnant women and their partners in making informed reproductive choices. On the other side, there are interests – particularly of the disability community and of future children – in equality, fairness and inclusion. The question is: how are regulators to ‘align the responsibilities that [they have] to support women to make informed reproductive choices about their pregnancies, with the responsibilities that [they have] . . . to promote equality, inclusion and fair treatment for all’?⁴ In response to which, the Council, being particularly mindful of the interests of future children – in an open future – and the interest in a wider societal environment that is fair and inclusive, recommends that a relatively restrictive approach should be taken to the use of NIPT.

In support of the Council’s approach and its recommendation, there is a good deal that can be said. For example, the Council consulted widely before drawing up the inventory of interests to be considered: it engaged with the arguments rationally and in good faith; where appropriate, its thinking was evidence-based; and its recommendation is not manifestly unreasonable. If we were to imagine a judicial review of the Council’s recommendation, it would surely survive the challenge.

However, if the Council had given greater weight to the interest in reproductive autonomy together with the argument that women have ‘a right to know’ and that healthcare practitioners have an interest in doing the best that they can for their patients,⁵ leading to a much less restrictive recommendation, we could say exactly the same things in its support.

In other words, so long as the Council – and, similarly, any regulatory body – consults widely and deliberates rationally, and so long as its recommendations are not manifestly unreasonable, we can treat its preferred accommodation of interests as acceptable. Yet, in such balancing deliberations, it is not clear where the onus of justification lies or what the burden of justification

² Nuffield Council on Bioethics, ‘Non-invasive Prenatal Testing: Ethical Issues’, (March 2017); for discussion, see R. Brownsword and J. Wale, ‘Testing Times Ahead: Non-Invasive Prenatal Testing and the Kind of Community that We Want to Be’, (2018) *Modern Law Review*, 81(4), 646–672.

³ Nuffield Council on Bioethics, ‘Genome Editing and Human Reproduction: Social and Ethical Issues’, (July 2018).

⁴ Nuffield Council on Bioethics, ‘Non-Invasive Prenatal Testing’, para 5.20.

⁵ Compare N. J. Wald et al., ‘Response to Walker’, (2018) *Genetics in Medicine*, 20(10), 1295; and in Canada, see the second phase of the Pegasus project, Pegasus, ‘About the Project’, www.pegasus-pegase.ca/pegasus/about-the-project/.

is; and, in the final analysis, we cannot say why the particular restrictive position that the Council takes is more or less acceptable than a less restrictive position.

Turning to the Council's second report, it hardly needs to be said that the development of precision gene-editing techniques, notably CRISPR-Cas9, has given rise to considerable debate.⁶ Addressing the ethics of gene editing and human reproduction, the Council adopted a similar approach to that in its report on NIPT. Following extensive consultation – and, in this case, an earlier, more general, report⁷ – there is a careful consideration of a range of legitimate interests, following which a relatively conservative position is taken. Once again, although the position taken is not manifestly unreasonable, it is not entirely clear why this particular position is taken.

Yet, in this second report, there is a sense that something more than balancing might be at stake.⁸ For example, the Council contemplates the possibility that genome editing might inadvertently lead to the extinction of the human species – or, conversely, that genome editing might be the salvation of humans who have catastrophically compromised the conditions for their existence. In these short reflections about the interests of 'humanity', we can detect a bigger picture of regulatory responsibilities.

27.3 THE BIGGER PICTURE OF REGULATORY RESPONSIBILITIES

In this part of the chapter, I sketch what I see as the bigger – three-tier – picture of regulatory responsibilities and then speak briefly to the first two tiers.

27.3.1 *The Bigger Picture*

My claim is that regulators have a first-tier 'stewardship' responsibility for maintaining the pre-conditions for any kind of human social community ('the commons'). At the second tier, regulators have a responsibility to respect the fundamental values of a particular human community, that is to say, the values that give that community its particular identity. At the third tier, regulators have a responsibility to seek out an acceptable balance of legitimate interests. The responsibilities at the first tier are cosmopolitan and non-negotiable. The responsibilities at the second and third tiers are contingent, depending on the fundamental values and the interests recognised in each particular community. Conflicts between commons-related interests, community values and individual or group interests are to be resolved by reference to the lexical ordering of the tiers: responsibilities in a higher tier always outrank those in a lower tier. Granted, this does not resolve all issues about trade-offs and compromises because we still have to handle horizontal conflicts *within* a particular tier. But, by identifying the tiers of responsibility, we take an important step towards giving some structure to the bigger picture.

27.3.2 *First-Tier Responsibilities*

Regulatory responsibilities start with the existence conditions that support the particular biological needs of humans. Beyond this, however, as *agents*, humans characteristically have the capacity to pursue various projects and plans whether as individuals, in partnerships, in groups, or in whole communities. Sometimes, the various projects and plans that they pursue will be

⁶ See, e.g., J. Harris and D. R. Lawrence, 'New Technologies, Old Attitudes, and Legislative Rigidity' in R. Brownsword et al. (eds) *Oxford Handbook of Law, Regulation and Technology* (Oxford University Press, 2017), pp. 915–928.

⁷ Nuffield Council on Bioethics, 'Genome Editing: An Ethical Review', (September 2016).

⁸ Nuffield Council on Bioethics, 'Genome Editing and Human Reproduction', paras 3.72–3.78.

harmonious; but often – as when the acceptability of the automation of healthcare and research is at issue – human agents will find themselves in conflict with one another. Accordingly, regulators also have a responsibility to maintain the conditions – conditions that are entirely neutral between the particular plans and projects that agents individually favour – that constitute the context for agency itself.

Building on this analysis, the claim is that the paramount responsibility for regulators is to protect, preserve, and promote:

- the essential conditions for *human* existence (given *human* biological needs);
- the generic conditions for human *agency* and self-development; and,
- the essential conditions for the development and practice of moral *agency*.

These, it bears repeating, are imperatives in all regulatory spaces, whether international or national, public or private. Of course, determining the nature of these conditions will not be a mechanical process. Nevertheless, let me indicate how the distinctive contribution of each segment of the commons might be elaborated.

In the first instance, regulators should take steps to maintain the natural ecosystem for human life.⁹ At minimum, this entails that the physical well-being of humans must be secured: humans need oxygen, they need food and water, they need shelter, they need protection against contagious diseases, if they are sick they need whatever treatment is available, and they need to be protected against assaults by other humans or non-human beings. When the Nuffield Council on Bioethics discusses catastrophic modifications to the human genome or to the ecosystem, it is this segment of the commons that is at issue.

Second, the conditions for meaningful self-development and agency need to be constructed: there needs to be sufficient trust and confidence in one's fellow agents, together with sufficient predictability to plan, so as to operate in a way that is interactive and purposeful rather than merely defensive. Let me suggest that the distinctive capacities of prospective agents include being able: to form a sense of what is in one's own *self*-interest; to choose one's own ends, goals, purposes and so on ('to do one's own thing'); and to form a sense of one's own identity ('to be one's own person').

Third, the commons must secure the conditions for an aspirant moral community, whether the particular community is guided by teleological or deontological standards, by rights or by duties, by communitarian or liberal or libertarian values, by virtue ethics, and so on. The generic context for moral community is impartial between competing moral visions, values, and ideals; but it must be conducive to 'moral' development and 'moral' agency in the sense of forming a view about what is the 'right thing' to do relative to the interests of both oneself and others.

On this analysis, each human agent is a stakeholder in the commons where this represents the essential conditions for human existence together with the generic conditions of both self-regarding and other-regarding agency. While respect for the commons' conditions is binding on all human agents, it should be emphasised that these conditions do not rule out the possibility of prudential or moral pluralism. Rather, the commons represents the pre-conditions for both individual self-development and community debate, giving each agent the opportunity to develop his or her own view of what is prudent, as well as what should be morally prohibited, permitted or required.

⁹ Compare, J. Rockström et al., 'Planetary Boundaries: Exploring the Safe Operating Space for Humanity' (2009) *Ecology and Society*, 14(2); K. Raworth, *Doughnut Economics* (Random House Business Books, 2017), pp. 43–53.

27.3.3 *Second-Tier Responsibilities*

Beyond the stewardship responsibilities, regulators are also responsible for ensuring that the fundamental values of their particular community are respected. Just as each individual human agent has the capacity to develop their own distinctive identity, the same is true if we scale this up to communities of human agents. There are common needs and interests but also distinctive identities.

In the particular case of the United Kingdom: although there is not a general commitment to the value of social solidarity, arguably this is actually the value that underpins the NHS. Accordingly, if it were proposed that access to NHS patient data – data, as Philip Aldrick has put it, that is ‘a treasure trove . . . for developers of next-generation medical devices’¹⁰ – should be part of a transatlantic trade deal, there would surely be an uproar because this would be seen as betraying the kind of healthcare community that we think we are.

More generally, many nation states have expressed their fundamental (constitutional) values in terms of respect for human rights and human dignity.¹¹ These values clearly intersect with the commons’ conditions and there is much to debate about the nature of this relationship and the extent of any overlap – for example, if we understand the root idea of human dignity in terms of humans having the capacity freely to do the right thing for the right reason,¹² then human dignity reaches directly to the commons’ conditions for moral agency.¹³ However, those nation states that articulate their particular identities by reference to their commitment to respect for human dignity are far from homogeneous. Whereas in some communities, the emphasis of human dignity is on individual empowerment and autonomy, in others it is on constraints relating to the sanctity, non-commercialisation, non-commodification and non-instrumentalisation of human life.¹⁴ These differences in emphasis mean that communities articulate in very different ways on a range of beginning-of-life and end-of-life questions as well as on questions of acceptable health-related research, and so on.

Given the conspicuous interest of today’s regulators in exploring technological solutions, an increasingly important question will be whether, and if so, how far, a community sees itself as distinguished by its commitment to regulation by rule and by human agents. In some smaller-scale communities or self-regulating groups, there might be resistance to a technocratic approach because automated compliance compromises the context for trust and for responsibility. Or, again, a community might prefer to stick with regulation by rules and by human agents because it is worried that with a more technocratic approach, there might be both reduced public participation in the regulatory enterprise and a loss of flexibility in the application of technological measures.

If a community decides that it is generally happy with an approach that relies on technological measures rather than rules, it then has to decide whether it is also happy for humans to be

¹⁰ P. Aldrick, ‘Make No Mistake, One Way or Another NHS Data Is on the Table in America Trade Talks’, *The Times*, (8 June 2019), 51.

¹¹ See R. Brownsword, ‘Human Dignity from a Legal Perspective’ in M. Duwell et al. (eds), *Cambridge Handbook of Human Dignity* (Cambridge University Press, 2014), pp. 1–22.

¹² For such a view, see R. Brownsword, ‘Human Dignity, Human Rights, and Simply Trying to Do the Right Thing’ in C. McCrudden (ed), *Understanding Human Dignity – Proceedings of the British Academy 192* (The British Academy and Oxford University Press, 2013), pp. 345–358.

¹³ See R. Brownsword, ‘From Erewhon to Alpha Go: For the Sake of Human Dignity Should We Destroy the Machines?’, (2017) *Law, Innovation and Technology*, 9(1), 117–153.

¹⁴ See D. Beylveled and R. Brownsword, *Human Dignity in Bioethics and Biolaw* (Oxford University Press, 2001); R. Brownsword, *Rights, Regulation and the Technological Revolution* (Oxford University Press, 2008).

out of the loop. Furthermore, once a community is asking itself such questions, it will need to clarify its understanding of the relationship between humans and robots – in particular, whether it treats robots as having moral status, or legal personality, and the like.

These are questions that each community must answer in its own way. The answers given speak to the kind of community that a group aspires to be. That said, it is, of course, essential that the fundamental values to which a particular community commits itself are consistent with (or cohere with) the commons' conditions.

27.4 AUTOMATED HEALTHCARE AND THE BIGGER PICTURE OF REGULATORY RESPONSIBILITY

One of the features of the NHS *Long Term Plan*¹⁵ – in which the NHS is described as ‘a hotbed of innovation and technological revolution in clinical practice’¹⁶ – is the anticipated role to be played by technology in ‘helping clinicians use the full range of their skills, reducing bureaucracy, stimulating research and enabling service transformation’.¹⁷ Moreover, speaking about the newly created unit, NHSX (a new joint organisation for digital, data and technology), the Health Secretary, Matt Hancock, said that this was ‘just the beginning of the tech revolution, building on our Long Term Plan to create a predictive, preventative and unrivalled NHS’.¹⁸

In this context, what should we make of the regulatory challenge presented by smart machines and devices that incorporate the latest AI and machine learning algorithms for healthcare and research purposes? Typically, these technologies need data on which to train and to improve their performance. While the consensus is that the collection and use of personal data needs governance and that big datasets (interrogated by state of the art algorithmic tools) need it a fortiori, there is no agreement as to what might be the appropriate terms and conditions for the collection, processing and use of personal data or how to govern these matters.¹⁹

In its recent final report on Ethics Guidelines for Trustworthy AI,²⁰ the European Commission (EC) independent high-level expert group on artificial intelligence takes it as axiomatic that the development and use of AI should be ‘human-centric’. To this end, the group highlights four key principles for the governance of AI, namely: respect for human autonomy, prevention of harm, fairness and explicability. Where tensions arise between these principles, then they should be dealt with by ‘methods of accountable deliberation’ involving ‘reasoned, evidence-based reflection rather than intuition or random discretion’.²¹ Nevertheless, it is emphasised that there might be cases where ‘no ethically acceptable trade-offs can be identified. Certain fundamental rights and correlated principles are absolute and *cannot be subject to a balancing exercise* (e.g. human dignity)’.²²

¹⁵ NHS, ‘NHS Long Term Plan’, (January 2019), www.longtermplan.nhs.uk.

¹⁶ *Ibid.*, 91.

¹⁷ *Ibid.*

¹⁸ Department of Health and Social Care, ‘NHSX: New Joint Organisation for Digital, Data and Technology’, (19 February 2019), www.gov.uk/government/news/nhsx-new-joint-organisation-for-digital-data-and-technology.

¹⁹ Generally, see R. Brownsword, ‘Law, Technology and Society’, Ch. 12; D. Schönberger, ‘Artificial Intelligence in Healthcare: A Critical Analysis of the Legal and Ethical Implications’, (2019) *International Journal of Law and Information Technology*, 27(2), 171–203.

For the much-debated collaboration between the Royal Free London NHS Foundation Trust and Google DeepMind, see, J. Powles, ‘Google DeepMind and healthcare in an age of algorithms’, (2017) *Health and Technology*, 7(4), 351–367.

²⁰ European Commission, ‘Ethics Guidelines for Trustworthy AI’, (8 April 2019).

²¹ *Ibid.*, 13.

²² *Ibid.*, emphasis added.

In line with this analysis, my position is that while there might be many cases where simple balancing is appropriate, there are some considerations that should never be put into a simple balance. The group mentions human rights and human dignity. I agree. Where a community treats human rights and human dignity as its constitutive principles or values, they act – in Ronald Dworkin’s evocative terms – as ‘trumps’.²³ Beyond that, the interest of humanity in the commons should be treated as even more foundational (so to speak, as a super-trump).

It follows that the first question for regulators is whether new AI technologies for healthcare and research present any threat to the existence conditions for humans, to the generic conditions for self-development, and to the context for moral development. It is only once this question has been answered that we get to the question of compatibility with the community’s particular constitutive values, and, then, after that, to a balancing judgment. If governance is to be ‘human-centric’, it is not enough that no individual human is exposed to an unacceptable risk or is not actually harmed. To be fully human-centric, technologies must be designed to respect both the commons and the constitutive values of particular human communities.

Guided by these regulatory imperatives, we can offer some short reflections on the three elements of the commons and how they might be compromised by the automation of research and healthcare.

27.4.1 *The Existence Conditions*

Famously, Stephen Hawking remarked that ‘the advent of super-intelligent AI would be either the best or the worst thing ever to happen to humanity’.²⁴ As the best thing, AI would contribute to ‘[the eradication of] disease and poverty’²⁵ as well as ‘[helping to] reverse paralysis in people with spinal-cord injuries’.²⁶ However, on the downside, some might fear that in our quest for greater safety and well-being, we will develop and embed ever more intelligent devices to the point that there is a risk of the extinction of humans – or, if not that, then a risk of humanity surviving ‘in some highly suboptimal state or in which a large portion of our potential for desirable development is irreversibly squandered’.²⁷ If this concern is well-founded, then communities will need to be extremely careful about how far and how fast they go with intelligent devices.

Of course, this is not specifically a concern about the use of smart machines in the hospital or in the research facility: the concern about the existential threat posed to humans by smart machines arises across the board; and, indeed, concerns about existential threats are provoked by a range of emerging technologies.²⁸ In such circumstances, a regulatory policy of precaution and zero risk is indicated; and while stewardship might mean that the development and application of some technologies that we value has to be restricted, this is better than finding that they have compromised the very conditions on which the enjoyment of such technologies is predicated.

²³ R. Dworkin, *Taking Rights Seriously*, revised edition (London: Duckworth, 1978).

²⁴ S. Hawking, *Brief Answers to the Big Questions* (London: John Murray, 2018) p. 188.

²⁵ *Ibid.*, p. 189.

²⁶ *Ibid.*, p. 194.

²⁷ See, N. Bostrom, *Superintelligence* (Oxford University Press, 2014), p. 281 (note 1); M. Ford, *The Rise of the Robots* (London: Oneworld, 2015), Ch. 9.

²⁸ For an indication of the range and breadth of this concern, see e.g. ‘Resources on Existential Risk’, (2015), [www.futureofflife.org/data/documents/Existential%20Risk%20Resources%20\(2015-08-24\).pdf](http://www.futureofflife.org/data/documents/Existential%20Risk%20Resources%20(2015-08-24).pdf).

27.4.2 *The Conditions for Self-Development and Agency*

The developers of smart devices are hungry for data: data from patients, data from research participants, data from the general public. This raises concerns about privacy and data protection. While it is widely accepted that our privacy interests – in a broad sense – are ‘contextual’,²⁹ it is important to understand not just that ‘there are contexts and contexts’ but that there is a Context in which we all have a common interest. What most urgently needs to be clarified is whether any interests that we have in privacy and data protection touch and concern the essential conditions (the Context).

If, on analysis, we judge that privacy reaches through to the interests that agents necessarily have in the commons’ conditions – particularly in the conditions for self-development and agency – it is neither rational nor reasonable for agents, individually or collectively, to authorise acts that compromise these conditions (unless they do so in order to protect some more important condition of the commons). As Bert-Jaap Koops has so clearly expressed it, privacy has an ‘infrastructural character’, ‘having privacy spaces is an important presupposition for autonomy [and] self-development’.³⁰ Without such spaces, there is no opportunity to be oneself.³¹ On this reading, privacy is not so much a matter of protecting goods – informational or spatial – in which one has a personal interest, but protecting infrastructural goods in which there is either a common interest (engaging first-tier responsibilities) or a distinctive community interest (engaging second-tier responsibilities).

By contrast, if privacy – and, likewise, data protection – is simply a legitimate informational interest that has to be weighed in an all things considered balance of interests, then we should recognise that what each community will recognise as a privacy interest and as an acceptable balance of interests might well change over time. To this extent, our reasonable expectations of privacy might be both ‘contextual’ and contingent on social practices.

27.4.3 *The Conditions for Moral Development and Moral Agency*

As I have indicated, I take it that the fundamental aspiration of *any* moral community is that regulators and regulatees alike should try to do the right thing. However, this presupposes a process of moral reflection and then action that accords with one’s moral judgment. In this way, agents exercise judgment in trying to do the right thing and they do what they do for the right reason in the sense that they act in accordance with their moral judgment. Accordingly, if automated research and healthcare relieves researchers and clinicians from their moral responsibilities, even though well intended, this might result in a significant compromising of their dignity, qua the conditions for moral agency.³²

Equally, if robots or other smart machines are used for healthcare and research purposes, some patients and participants might feel that this compromises their ‘dignity’ – robots might not

²⁹ See, for example, D. J. Solove, *Understanding Privacy* (Cambridge, MA: Harvard University Press, 2008); H. Nissenbaum, *Privacy in Context* (Palo Alto, CA: Stanford University Press, 2010).

³⁰ B. Koops, ‘Privacy Spaces’, (2018) *West Virginia Law Review*, 121(2), 611–665, 621.

³¹ Compare, too, M. Brincker, ‘Privacy in Public and the Contextual Conditions of Agency’ in T. Timan, et al. (eds), *Privacy in Public Space* (Cheltenham: Edward Elgar, 2017), pp. 64–90; M. Hu, ‘Orwell’s 1984 and a Fourth Amendment Cybersurveillance Nonintrusion Test’, (2017) *Washington Law Review*, 92(4), 1819–1904, 1903–1904.

³² Compare K. Yeung and M. Dixon-Woods, ‘Design-Based Regulation and Patient Safety: A Regulatory Studies Perspective’, (2010) *Social Science and Medicine*, 71(3), 502–509.

physically harm humans, but even caring machines, so to speak, ‘do not really care’.³³ The question then is whether regulators should treat the interests of such persons as a matter of individual interest to be balanced against the legitimate interests of others, or as concerns about dignity that speak to matters of either (first-tier) common or (second-tier) community interest.

In this regard, consider the case of Ernest Quintana whose family were shocked to find that, at a particular Californian hospital, a ‘robot’ displaying a doctor on a screen was used to tell Ernest that the medical team could do no more for him and that he would soon die.³⁴ What should we make of this? Should we read the family’s shock as simply expressing a preference for the human touch or as going deeper to the community’s constitutive values or even to the commons’ conditions? Depending on how this question is answered, regulators will know whether a simple balance of interests is appropriate.

27.5 CONCLUSION

In this chapter, I have argued that it is not always appropriate to respond to new technologies for healthcare and research simply by enjoining regulators to seek out an acceptable balance of interests. My point is not that we should eschew either the balancing approach or the idea of ‘acceptability’ but that regulators should respond in a way that is sensitised to the full range of their responsibilities.

To the simple balancing approach, with its broad margin for ‘acceptable’ accommodation, we must add the regulatory responsibility to be responsive to the red lines and basic values that are distinctive of the particular community. Any claimed interest or proposed accommodation of interests that crosses these red lines or that is incompatible with the community’s basic values is ‘unacceptable’ – but this is for a different reason to that which applies where a simple balancing calculation is undertaken.

Most fundamentally, however, regulators have a stewardship responsibility in relation to the anterior conditions for *humans* to exist and for them to function as a *community of agents*. We should certainly say that any claimed interest or proposed accommodation of interests that is incompatible with the maintenance of these conditions is totally ‘unacceptable’ – but it is more than that. Unlike the red lines or basic values to which a particular community commits itself – red lines and basic values that may legitimately vary from one community to another – the commons’ conditions are not contingent or negotiable. For human agents to compromise the conditions upon which human existence and agency is itself predicated is simply unthinkable.

Finally, it should be said that my sketch of the regulatory responsibilities is incomplete – in particular, concepts such as the ‘public interest’ and the ‘public good’ need to be located within this bigger picture; and, there is more to be said about the handling of horizontal conflicts and tensions *within* a particular tier. Nevertheless, the ‘take home message’ is clear. Quite simply: while automated healthcare and research might be efficient and productive, new technologies should not present unacceptable risks to the legitimate interests of humans; beyond mere balancing, new technologies should be compatible with the fundamental values of particular communities; and, above all, these technologies should do no harm to the commons’ conditions – supporting human existence and agency – on which we all rely and which we undervalue at our peril.

³³ Compare R. Brownsword, ‘Regulating Patient Safety: Is It Time for a Technological Response?’, (2014) *Law, Innovation and Technology*, 6(1), 1–29.

³⁴ See M. Cook, ‘Bedside Manner 101: How to Deliver Very Bad News’, *Bioedge* (17 March 2019), www.bioedge.org/bioethics/bedside-manner-101-how-to-deliver-very-bad-news/12998.