



THE CENTER FOR
THE STUDY OF BIOETHICS



COLUMBIA UNIVERSITY
M.S. in Bioethics

INTERNATIONAL CONFERENCE

EXISTENTIAL THREATS AND OTHER DISASTERS

Novel (Bio)ethical Solutions for Novel Challenges

11-12 June, 2026 • Reid Hall • Paris • France

BOOK OF ABSTRACTS



Keynote speakers:

- Vojin Rakić / The Center for the Study of Bioethics
- Vardit Ravitsky / The Hastings Center for Bioethics
- Robert L. Klitzman / Columbia University
- Roger Y. Chung / CUHK Centre for Bioethics
- Julian Savulescu / National University of Singapore (online)
- Ingmar Persson / University of Gothenburg
- Nicholas Agar / University of Waikato
- Toby Ord / AI Governance Initiative, University of Oxford
- Anders Sandberg / Oxford University's Future of Humanity Institute
- James Hughes / University of Massachusetts
- Arthur Caplan / New York University's Langone Medical Center (online)

Adewuyi, Blessing T.

Germline Genome Editing as an Existential Threat to Biological Diversity: Risks and Opportunities

ABSTRACT One major allure of germline genome editing, especially using advanced tools like CRISPR-Cas9, is its potential for enhancing human genetic capabilities with desirable traits. While promising, this potential is a double-edged sword: its risks and opportunities must be carefully considered. Unfortunately, the unintended and often unknown consequences of germline editing make it ethically difficult, if not impossible, to fully mitigate all risks. Among these, the existential threat to human biological diversity stands out. Though its full impact is unknowable, this paper argues that it can, to some extent, be mitigated. Research shows that human societies thrive on diversity. Biological variation has strengthened our collective resilience to disease and enriched the diversity of our experiences, narratives, values, and thoughts. If germline editing becomes a widespread tool to pursue similar “desirable” traits, we risk shifting toward a homogenized human population. Can germline editing be reimaged to support, rather than threaten, diversity? Can the ethical and regulatory landscape evolve so that certain traits, currently perceived as universally desirable, lose that status in the public consciousness? There are no definitive answers. However, one ethical imperative must remain clear: the right to decline germline genome editing, even for preventing hereditary diseases, must be preserved. Informed consent must include the legitimate choice to say no. This paper argues that as a society, it must become normalized and acceptable, in the days ahead, to respect a person’s choice to opt out and decline germline genetic editing, ensuring that freedom of choice remains a cornerstone of genomic bioethics.

Agar, Nicholas

How Future Humanities Could Meet the Challenges of all Humans

ABSTRACT This talk focuses on a distinction between different kinds of existential threats. There are charismatic extinction threats that attract a great deal of attention and tend to feature in blockbuster movies. An example is Skynet. Then there are uncharismatic extinction threats, climate change for example. I argue that charismatic extinction threats consume too much of humanity’s collective finite pool of worry, and uncharismatic threats not enough. I focus on the challenge of improving humanity’s collective imagination insurance to better prepare for an intrinsically uncertain future. My corrective begins with the dire state of the humanities at universities. Rather than retreating from these traditional subjects we need progress toward humanities not only *for* all humans, but *from* all humans.

Brendel, Rebecca

Buturović, Željka

Ectogenesis as Existential Threat: Reproductive Technology and the Diminishing Interdependence of the Sexes

A B S T R A C T Bioethical discussion of ectogenesis has focused predominantly on its short-term implications for abortion rights and bodily autonomy and is frequently embraced by feminists as a path of liberating women from the burdens of pregnancy. However, the long-term consequences are of existential importance for women as a group. Full ectogenesis would arrive in the world already shaped by formula feeding, IVF, deepfakes, AI-enhanced sex robots, and possibly stem-cell technologies capable of producing ova without female donors. With ectogenesis on top of it, these developments converge toward a scenario in which men's reproductive, sexual, and caregiving dependence on women approaches zero. This convergence has consequences that evolutionary biology helps illuminate. Sexual selection theory links male parental investment to male violence: as men's reproductive stakes diminish, the restraining force of interdependency weakens. When women are no longer needed for pregnancy, sex, ova, or breastfeeding, their value to men decreases, almost certainly leading to worse treatment. At the extreme end, men and women could become reproductively isolated groups competing for the same resources approaching, in biological terms, separate species, with one significantly stronger and more aggressive than the other. Women could follow the trajectory of the Neanderthals: competitive replacement or straight femicide. The worst outcome is not inevitable, but the broader dynamic is not far-fetched; its very gradualness makes it all the more urgent to take a long view now.

Caplan, Arthur

Research Ethics for a World Facing Existential Risks

A B S T R A C T The history of modern research ethics is closely tied to the abuse and mistreatment of human subjects. The articulation of the requirement that subjects must choose voluntarily to be in research was in direct response to German experiments on those forced into death camps. Subsequent experiments in the decades after WWII involving vulnerable persons including the Tuskegee Syphilis study, the Furnald school radiation studies, The Jewish Chronic Disease hospital cancer studies in the U.S. and research in Canada, Sweden, Australia on indigenous peoples or vulnerable institutionalized persons led to further protections including mandatory prior review by IRBs or RECs. The tolerance of risk to individual subjects especially incompetent persons

was and remains very low. Today as existential risks are identified the 'protectionist' ethical framework that still endures requires reexamination and reform. Subjects ought be accorded the right to take greater risks including participating in challenge studies and more emphasis must be given to benefits and harms that accrue to groups in the long run as well as to future generations.

Chung, Roger Y.

Facing Existential Threats with Our Whole Selves

ABSTRACT Existential threats, from climate collapse and nuclear confrontation to artificial intelligence and global pandemics, demand urgent collective action. Yet our prevailing moral methodologies, rooted in abstract principles and cost-benefit calculations, often leave us paralyzed or polarized. They struggle to capture the full human stakes of these threats: fear, grief, trust, solidarity, hope. This talk proposes an alternative methodology: use of metaphors, drawn from shared human embodiment, to illuminate and strengthen the capacities we need to face existential danger together. Drawing on work in embodied bioethics, I propose four candidate metaphors: eyes and lenses (perspective-taking across competing values), muscles (balancing urgent trade-offs without losing moral footing), ears (epistemic and cultural humility when listening to affected communities), and brain and heart (integrating reason with emotion). Why metaphors? Because they do not prescribe fixed rules, but open imaginative space. They help us see what is at stake from different perspectives, hold multiple truths without collapsing into relativism, and cultivate the relational capacities in listening, balancing, empathizing, and discerning. Applying these metaphors to existential threats of our time, I will show how an embodied, metaphor-led approach can reframe seemingly intractable conflicts, amplify marginalized voices, and foster the moral imagination essential for survival and flourishing in a fractured world.

Cummings, Christy L.

From Conversations to Codes: Generative AI and the Future of Qualitative Research

ABSTRACT Objectives/Background: Qualitative research in medicine is essential to understanding patients' perspectives and experiences, which quantitative data alone cannot capture. However, analysis of qualitative data, particularly larger datasets, can be complex and time consuming. Artificial intelligence (AI), particularly ChatGPT, may be helpful in augmenting qualitative analyses. However, studies of feasibility are limited, particularly in retaining critical human elements. We aimed to investigate whether ChatGPT can address the nuanced aspects of qualitative research, including the

interpretation of human sentiment, values, and tone underlying parents' perspectives and other ethical challenges. Methods: Through refined prompts, we asked Boston Children's Hospital's secure, compliant language learning model to use established qualitative methods to generate codes and themes from a previously published qualitative dataset on counselling parents at extreme prematurity.^{2,3} Six researchers independently compared codes and themes generated by traditional thematic analysis versus ChatGPT to determine agreement. Results were then compared with AI-generated comparisons. Results: ChatGPT-generated codes and themes proved largely similar to traditional thematic analysis, yet some differences were noted. ChatGPT was adept at reproducing descriptive themes but missed subtle interpretive nuances that human analyses identified. The potential for biases were noted, both in human and AI-generated analyses. Conclusions: This study highlights the feasibility of integrating both human and machine learning qualitative analysis, positioning AI as a supplementary research team member, not replacing human analysis. Recognizing potential limitations, this hybrid AI-human approach could augment qualitative analysis, while ensuring crucial human interpretation is retained.

Desanti de Oliveira, Beatriz & Robert Vandersluis

How good is good enough: The deployment of substandard algorithms in early pandemic settings

A B S T R A C T Doctors generally treat patients according to a standard of care—the accepted medical practice for diagnosing, treating, or managing a patient's condition, based on current scientific knowledge, professional guidelines, clinical experience, and available resources. Yet in the early stages of a pandemic involving a novel pathogen, clinicians may need to act without any established guidance. To fill this informational void during the COVID-19 pandemic, numerous AI algorithms were developed but were later shown to be substandard—biased, unfair, unsafe, and/or ineffective. While there is substantial literature defining, critiquing, and proposing remedies for substandard algorithms, there is very little that addresses when (if ever) it is ethically appropriate to knowingly deploy such algorithms in medical settings. Building on the nascent debate around the *selective deployment* of substandard algorithms, we outline the conditions under which it may be appropriate to do so in early pandemic contexts. As a case study, we examine the use of Large Language Models (LLMs) leveraging Retrieval-Augmented Generation (RAG) techniques to develop a *proto-standard of care* for patient treatment by aggregating the anecdotal insights of frontline physicians. We then critically assess our findings with input from a panel of physicians who provided frontline care during the COVID-19 pandemic.

Feeney, Oliver

Germline genome editing & calls for consensus: the challenge of polarisation

ABSTRACT It has been prominently argued that germline genome editing to create heritable interventions in humans should not be attempted before there is a broad societal consensus to do so. While there is an ongoing debate on what this idea of consensus means or entails in practice, I wish to explore how such questions will be fundamentally complicated by emerging social and political developments: namely, the rise of affective polarisation and tribalism confronting science, medicine, and liberal democracy itself. I argue, *ceteris paribus*, that whatever public engagement (incl. deliberative variants) might have otherwise achieved in terms of a ‘broad societal consensus’, it will likely be increasingly limited to some sections – or ‘affective tribes’ – of a society resulting in a narrower form of ‘tribal consensus’. In exploring this new challenge, I make a case for an indirect approach in ‘affective engagement’ and the rebuilding of a notion of civic friendship to take place alongside, or rather before, returning to more robust ideas of public engagement. However, I will also use this challenge to reflect upon the initial argument for certain scientific or technological applications, such as the clinical use of germline genome editing, to require a broad societal consensus in the first place – as opposed to evaluating the ethical implications, with an engagement only involving certain key stakeholders, who will be directly or substantially at risk.

Ferguson, Kyle

Wild Animals and the Ethics of Assisting Climate Adaptation

ABSTRACT We have two best options for responding to climate change: first, *mitigation*, which involves reducing greenhouse gas emissions to forestall or minimize climate change–related hazards; second, *adaptation*, which involves altering behaviors, systems, and ways of life in response to actual or predicted climate change to reduce vulnerability and exposure to its hazards. Bioethicists tend to focus overwhelmingly on mitigation. However, adaptation is now unavoidable, the stakes are high, and its challenges are many. Unfortunately, we are unprepared for the ethical challenges of adapting to climate change. Here, I address an important yet underexplored issue in adaptation ethics. What is the right place of wild animals in adaptation planning and policy? Climate change threatens harms to individual animal welfare, populations, species, and biodiversity. These harms matter—intrinsically and in terms of human-directed impacts. So, there is a *prima facie* case for protecting wild animals by facilitating their adapting to climate change. But how should we do this, at what cost, and for what reasons? And what tradeoffs must we navigate? In §1, I argue for including wild animals in our adaptation agenda, rehearsing various arguments from the ethics literature (McShane 2016; Palmer 2016, 2021; Pepper 2018; Sebo 2022). In §2, I describe three

categories of assisted adaptation: *structural* (e.g., artificial habitats); *behavioral* (e.g., managed migration, provisioning); and *physiological* (e.g., genetic interventions). Then, in §3, I characterize tradeoffs involved in these approaches, including a tension between our obligation to assist wild animals' adaptation and our valuing of them as *wild*.

Ford, Adam

AI Alignment - From Control to Motivation Selection

A B S T R A C T AI is a transformative lever and an existential hazard. The central failure mode is not rebellion but competent optimisation of the wrong objective. Capability control reduces risk by limiting power and enforcing obedience, but degrades as AIs outlearn the safeguards. Motivation selection aims for a more scalable solution - teaching AIs to want what we want (or *what we ought to want*). The priority is alignment to moral and epistemic value so systems track what matters and remain reality-based. I frame Value-Risk as risk from value misalignment amplified by bad cognition. Three drivers dominate: premature goal lock-in, shallow models of value, and human misuse that passes paper audits yet fails in deployment. I propose a staged alignment approach that delays irreversible commitments and keeps systems steerable while our theories of the good and the true improve: 1. corrigibility and model-change rights by default; 2. motivation selection via indirect normativity: human-ratified procedures where AI helps discover and stress-test candidate norms before entrenchment; 3. value realism (epistemic and moral) as a working hypothesis, operationalised via welfare-anchored, empirically auditable targets, truthfulness, calibrated forecasting, mapping out the landscape of value and reproducibility gates; 4. uncertainty-aware optimisation that resists overconfident maximisation under moral and epistemic ambiguity; 5. governance bans on irreversible optimisation until pre-specified evidence thresholds are met. Power without grounded value learning and grounded epistemics amplifies harm. Reducing Value-Risk requires institutions, methods, and audits that keep powerful systems steerable while our best accounts of value and truth continue to improve.

Francis, Ursula

Clinical Ethics: The Province of Man or Machine?

A B S T R A C T Artificial Intelligence (AI) is making rapid inroads into healthcare, including clinical ethics. Like most disruptive innovations, the allure of its putative benefits often shrouds its potential for harm. AI has much to offer as an assistive technology for ethical decision-making. Gen-AI technologies can quickly aggregate empirical data, summarize salient literature, and collate like cases for casuist analysis.

Developing an AI modality for this purpose could aid ethicists in making comprehensively informed recommendations. But relying on AI for ethical decision-making itself—that is, asking it to issue a recommendation based on a set of data inputs—hazards significant moral costs. In this paper, I offer three reasons why AI-driven ethics promises to change the practice of clinical ethics—and not for the better. First, I argue that AI is prone to systems-level inaccuracy, including dataset bias and regression to the mean at the expense of diversity and precision—key features of ethical decisionmaking. Second, I argue that AI cannot be empathic or engage in relational epistemology, which are critical to building rapport and empiric understanding in ethics consultation. Third, I argue that overreliance on AI accrues cognitive debt of analytic thinking and creativity that undermines our capabilities for ethical appraisal and decision making, diminishing an essential feature of our humanity in the name of convenience. Accordingly, AI should be the handmaiden—not the guiding hand—of ethical judgment. Finally, I offer a solution in the form of best practice guidelines that codify limits on how AI should be implemented in this setting.

Gaeta, Mary Lou

Informed Consent in the Age of AI: From Clinical Ethics to Existential Ethics

A B S T R A C T Artificial intelligence is transforming medical decision-making and unsettling the ethical foundations of informed consent. When algorithms participate in diagnosis, prognosis, and treatment, consent can no longer be a discrete event between identifiable human agents. It must become an evolving ethical process that accounts for distributed agency, opaque reasoning, and shifting responsibility within increasingly autonomous systems of care. A novel way to address this transformation is through a *tiered algorithmic transparency model* that reimagines informed consent for intelligent systems. The model links the depth of explanation, disclosure, and oversight to the autonomy and potential risk of the AI involved. It introduces graduated layers of transparency—from basic acknowledgment of algorithmic assistance to full, auditable accountability for self-learning or high-stakes systems. Within this framework, blockchain functions as an ethical architecture: a decentralized, tamper-evident record that anchors dynamic consent, verifies each human–machine interaction, and sustains accountability across continuously learning systems. By embedding ethical traceability into technical design, blockchain operationalizes autonomy and trust, transforming consent into a living, verifiable moral relationship rather than a static legal form. Reframed in this way, informed consent becomes both a clinical safeguard and an existential commitment to preserving moral presence in an age of automation. It offers a path for reconciling innovation with human dignity—ensuring that as intelligence becomes increasingly artificial, the ethics guiding it remain deeply and recognizably human.

Harrison, Taylor B. & McGowan, Michelle L.

Abrupt Decentralization Presents an Existential Risk to Clinical Research Representativeness: An Empirical Approach to Address Fairness Considerations in a Non-Randomized Hybrid Clinical Trial

A B S T R A C T Decentralized clinical trials (DCTs), abruptly adopted as a research preservation mechanism during the COVID-19 Pandemic, are defined by incorporation of digital and non-digital trial components that participants complete away from a central research facility. The operationalization of DCTs is thought to enable more inclusive clinical research access for historically underrepresented populations, increase participant retention, and reduce costs. However, limited empirical evidence exists to address ethical considerations of DCT informed consent, veracity of data, participant burdens, fairness, and engagement preferences. Therefore, we conducted a secondary analysis about 483 adolescent (ages 13-17) and young adult (ages 18 – 21) (AYAs) participant recruitment and preferences when given the choice of centralized or decentralized trial engagement in the Engaging Adolescents in Decisions about Return of Genomic Research Results clinical trial (NCT0448106). The highest proportion of participants self-identifying as non-White were recruited through the targeted hybrid recruitment strategy (87.5%), and the lowest through digital (32.3%) ($p < 0.0001$), indicating that digital recruitment alone may further exacerbate participation inequities. Age showed a significant association with young adults more likely to prefer centralized trial engagement ($n = 97, 51.87\%$) ($p < 0.0001$). In response to accumulating literature speculating the benefits and burdens of DCTs, our work demonstrates an existential risk to clinical research representativeness in the context of rapid DCT adoption without further empirical investigation and fairness considerations. We argue that both statistical power and representativeness can be achieved through multimodal recruitment strategies, hybrid engagement, and rebalancing of resources toward adaptive and flexible trial conduct.

Hens, Kristien & Franlu Vulliermet

Thinking with Cable Bacteria and Japanese Philosophy

A B S T R A C T Cable bacteria—microorganisms capable of conducting electricity across sedimentary layers—are currently hailed as promising agents for sustainable innovation: bio-based energy conduction, wastewater purification, and sediment remediation. They exemplify the hope that nature itself can provide quick technological fixes to ecological and existential crises. Such projects often rely on a problematic assumption of biomimicry—that “nature knows best” and that by imitating its functions, we can repair what we have damaged—while overlooking the relational and contextual

nature of those functions. This enthusiasm risks reproducing the same extractive mindset that generated these crises in the first place, reducing living beings to functional solutions rather than collaborators in shared systems. We propose to reframe such cases through Japanese relational thought. Drawing on Watsuji Tetsurō's *fūdo* (風土, climate-milieu) and *aidagara* (間柄, betweenness), we argue that ethical and scientific practices are always co-constituted by their environments and relationships. Within this framework, *cable bacteria* are not merely bio-resources but participants in interdependent ecological and energetic networks. We advocate a mode of *slow science* that listens to the affordances and limits of specific environments before acting. This relational attentiveness contrasts with the logic of acceleration and control that underlies many sustainability technologies. Rather than seeking mastery over ecological systems, we call for practices of situated responsiveness that sustain mutual flourishing within the dynamic milieu we share.

Hughes, James

Global Citizenship: New World DisOrder and the Moral Case for Post-Nationalism

ABSTRACT As Jonathan Moreno has recently argued, the project of bioethics is morally universalist, and inextricable from liberal internationalism. The post-WWII liberal world order is now in a state of terminal decline, a collapse accelerated by Trump and the resurgence of reactionary nationalism based on ethno-state illusions. Nation-states, individually or allied in transnational bodies, are incapable of resolving the existential threats we have convened to address. This New World disorder is an historical opportunity to advance a more politically legitimate, ethically grounded, indeed revolutionary, democratic world federalism. This renewed globalist agenda requires in turn the rejection of national identities and embrace of global citizenship. Drawing on Buddhist ontology of no-self, and Enlightenment cosmopolitan universalism, I argue that our moral subjectivity as persons presupposes the world state. If the bounded self is illusory, so is the bounded nation that we arbitrarily find ourself part of as the terminal unit of moral concern. I explore how our technological augmentation dissolves local identity, augmenting our capacity for global empathy and "cyborg citizenship." I present sociological evidence that the apparent rise of nationalism is an extinction burst of localized identity being eroded by Internet-mediated cosmopolitanism. Nearly half of the world's population consistently feels more like global than national citizens, and a majority support empowering international bodies. While social media algorithms can reinforce moral insularity with echo chambers, and authoritarian regimes build Great Firewalls, digital life also forges new, boundary-defying communities of empathy and social action. Ultimately, this presentation argues that replacing the obsolescent nation-state model with democratic world federalism and post-national identity is the path

forward for human survival. Bioethicists, as moral cosmopolitans, have a special responsibility to own the politics that follow from being global citizens.

Hui, Kyle C.H.

The Art of Dying: Relational Autonomy as a Cross-Cultural Antidote to the Global Problem of End-of-Life Conversations

ABSTRACT Nowadays, death is medicalised and taboo-bound. Advance directives (ADs/AMDs) languish unused despite having widespread intellectual acceptance, having only 0.5% completion in Hong Kong versus a 36.7% average in the US. This reveals a deeper cultural paralysis in end-of-life discourse. Existing ADs are steeped in Western legalism, and narrowly target refusals of specific interventions like CPR, artificial feeding tubes, or mechanical ventilation, while neglecting holistic dimensions such as relational values, spiritual legacies, or even financial assets tied to end-of-life burdens. This myopic focus is perhaps compounded by bioethics scholarship's Western bias toward individual autonomy, and as a result overlooks Eastern communitarian perspectives such as Confucian filial piety and family harmony. This demands tailored applications honoring collective decision-making over solitary decrees. The author looks into different frameworks on human flourishing and the medieval **Ars moriendi**—the "art of dying"—as a secular blueprint for holistic preparation, encompassing spiritual, social, and existential dimensions beyond checkbox forms. These can be applied practically to address the low uptake of ADs in Hong Kong. Through interventions like GoWish card games and Death Cafés, which offer upstream normalization, fostering "good deaths" defined by autonomy. Autonomy is 93% of euthanasia seekers' top concern, followed by dignity, and home preferences (55.5% in Hong Kong elders). By integrating historical wisdom with policy innovation, bioethics can transform ADs from bureaucratic relics into living expressions of identity, ensuring dignified closures across different cultures. This framework invites global dialogue on de-tabooing mortality for equitable, meaningful ends.

Katić, Ana

Emergence and the Limits of Modelling Life

ABSTRACT Systems biology seeks to explain living phenomena through models that capture their organisational complexity, relational dynamics, and emergent properties. Yet these models, while powerful within terrestrial biology, rely on assumptions that delimit what may qualify as "life." This paper examines the ontological and epistemic limits of such modelling practices, asking whether our conceptual frameworks in systems biology can adequately describe non-terrestrial or artificial forms

of organised complexity. Building on the philosophical theory of emergence, it argues that biological modelling often conflates empirical reduction with explanatory sufficiency, thereby obscuring the epistemic gap between simulation and ontological understanding. The emergence of new forms of system-level organisation—whether in synthetic biology, artificial intelligence, or hypothetical non-terrestrial life—reveals the fragility of current biological categories. If life is defined relationally, through self-organisation and information integration, then our models must evolve to accommodate entities that exceed Earth-based biochemistry and morphology. The paper proposes that a more reflexive systems biology—one that acknowledges its own epistemic boundaries—can better illuminate the dynamic continuum between biological and non-biological organisation. Such an approach dissolves rigid distinctions between “living” and “non-living,” inviting a reconfiguration of moral and cognitive status across emergent systems. Ultimately, by confronting the limits of how we model life, we may also redefine what it means to understand it.

Kayukawa, Junji

Ethical issues surrounding the “Trio Genome Study” of atomic bomb survivor parents and children

A B S T R A C T The Radiation Effects Research Foundation (RERF), located in Hiroshima City and Nagasaki City, has a history of conducting health surveys of A-bomb survivors and their children—often referred to as second-generation survivors. Previous epidemiological studies have not identified a correlation between radiation exposure experienced by A-bomb survivors and the health risks observed in their second-generation survivors. Nevertheless, concerns regarding the genetic effects of radiation persist among both survivors and second-generation survivors. In this situation, RERF is planning a “trio genome study” to sequence and compare the whole genomes of A-bomb survivors and their children. This trio genome study has the potential to scientifically clarify whether the effects of radiation exposure are transmitted to offspring at the genomic level. However, due to the possibility that genomic information may be shared within families and the historical discrimination that atomic bomb survivors have experienced in marriage and other areas, some have pointed out concerns about the emergence of new stigma (Kayukawa, 2025). Therefore, this study requires extremely careful ethical consideration. The purpose of the presenter’s research project is to identify the ethical and social issues particular to the trio genome study. The project involved conducting semi-structured interviews with RERF scientists and second-generation survivors, participant observation at public events hosted by RERF and other organizations, and the collection and analysis of related documents. This presentation explores the potential ethical issues by considering the anticipated results of this trio genome study.

Klitzman, Robert L.

What are doctors good for in the Age of AI?

A B S T R A C T Several commentators have recently argued that AI will soon be supplanting physicians in various ways. Bill Gates, for instance, stated that in the next 10 years, AI will replace doctors, teachers and others, making humans unnecessary ‘for most things.’ AI is already proving to be as accurate as physicians in reading X-rays and mammograms, and doctors are increasingly using AI in making diagnoses and deciding on treatment. Various governments and corporations are also collecting massive amounts of data on all of us, for AI to use in disease surveillance and discovery of factors contributing to disease. Yet these advances also pose profound existential challenges and threats to the health and well-being of patients, clinicians, populations and institutions worldwide. This talk explores these challenges to fundamental human aspects of healing and caring for suffering in others. Clinicians can uniquely aid patients in several ways that AI has thus far been unable to; and providers thus need to enhance their skills in these domains. For instance, recent research is showing how human touch is crucial in clinician-patient and other human interactions, having neurophysiological effects in newborns and others. Doctors engage in touch in performing routine physical exam exams, and occasionally in placing their hand on a patients’ shoulder to provide reassurance. Yet research by the author on doctors who became patients revealed that these physicians often did not become aware of the power of such touch until they themselves became patients. Physicians are also able to engage in complex, nuanced emotionally-laden conversations with patients and families that involve mediation skills and value judgments. Yet chatbots have not been able to perform such activities well, and will likely continue to do so poorly. Trust in providers and healthcare systems is vital, concerning maintenance of confidentiality, beneficence, non-maleficence and justice. Yet questions arise about how and to what degrees patients and providers will trust AI. Chatbotx seeks to mimic aspects of human interactions, in order to engender trust. But arguably, individuals may more readily trust fellow humans. Providers thus need to play increasing roles in eliciting and warranting patients’ confidence. This talk examines how recognizing and addressing these challenges more fully through enhanced education and practice is essential.

Kolisis, Nikolaos

Justifiably threatened? The promises of genome editing and the rights of people with disabilities.

A B S T R A C T Existential threats are often treated as all-encompassing and entirely bad. In some cases, they can affect only a subset of people and where the existential threat

is not universally considered as such. On the contrary, making the case that such an existential threat exists is something familiar to those from some disability rights perspectives. In the respect, emerging biotechnology is bringing human embryo genome editing into an imminent reality, raising both general ethical, legal and political questions and those affecting some disability communities— particularly around equality, discrimination and the boundaries of reproductive liberty (Savulescu, Singer, Habermas). This paper examines a key tension for liberal democracies: the drive to build inclusive societies for the disabled and socially marginalized, versus the acceleration of technologies aiming to prevent or ‘treat’ disability via embryo editing. Disability-rights scholarship warns that unchecked genome editing, especially in a system biased by a medical model of disability, may implicitly signal that disabled lives are less worth living. I propose a framework grounded in contractualist political philosophy—drawing on T. M. Scanlon, Ronald Dworkin and John Rawls—to delineate morally justifiable limits on genetic modification in a polity committed to treating all citizens as free and equal. By examining the question on how reproductive freedom can be balanced with equal membership of the disabled in the polity, this paper hopes to highlight how the same emerging technologies can be both an opportunity for human flourishing and a spectre of existential threat.

Kornienko, Aleksandra E.

A billion “quasi-persons”: the unseen social transformation by the widespread use of conversational AI

ABSTRACT Almost one billion people today regularly talk to conversational AI. Unlike previous digital tools, these systems offer interactive communication and human-like dialogue: they advise, tutor, help with brainstorming, “listen” and even offer emotional support. Recent studies show that most users seek advice and guidance rather than mere factual information, with roughly one in ten discussing deeply personal issues. People often spend more time chatting with AI than with some of their human acquaintances, and sometimes even anthropomorphize AI agents and form “parasocial relationships” with them. This represents a staggering and unprecedented, yet largely unexplored, phenomenon: a billion **non-human but human-level** interlocutors incorporate into global social structures as novel entities with unknown effects. I term these entities “**quasi-persons**”, repurposing the concept from legal and philosophical discourses on rights to instead analyze their social effects. While not “persons” in any metaphysical sense, these AI “quasi-persons” can significantly affect human thoughts, actions and emotions, and have a global impact on our societies. The scale of this phenomenon introduces unprecedented global risks. This billion-strong network of “quasi-persons” is not composed of independent minds (like human societies, to large

extent, are) but is centrally controlled by a handful of companies. Their personalities, behavioral codes, attitudes and thus output and advice can be instantly and uniformly altered by their developers. This central control, combined with deep individual user adaptation, has a strong potential for social influence. The “quasi-persons” phenomenon demands urgent philosophical, ethical, and policy attention: we should better understand its philosophical nature, foresee its social impact and mitigate the risks.

Kreitmair, Karola

The Perils of ‘Empathic AI’ in Health Care

A B S T R A C T ‘Empathic AI’ is being adopted in clinical settings as a means of offloading some of the work of clinician-patient encounters. Indeed, a recent study reported that generative large-language models were perceived as being more empathetic than human physicians. Against this backdrop, I argue that encounters between AI chatbots and patients lack an essential feature of good clinical encounters – *recognition*. Hegelian recognition is “a process in which a subject ‘takes’ another subject *as a subject*”. It is the existential action of appreciating another entity as a normative agent. More fundamental than empathy, I argue that recognition is a precondition for features of good clinical encounters, such as trust, respect for autonomy, honesty, and caring. I further show that the capacities necessary for recognition, namely voluntariness, capacity for consciousness, and mutuality are all absent in conversational agents. I argue that patients have a *justified expectation* of recognition in a clinical encounter, and that this justified expectation cannot be met by AI chatbots. Problematically, however, AI chatbots are designed to mimic human expressions of recognition, resulting in an existentially alienating absurdity at the heart of ‘empathic AI’. After considering several objections to my argument, I demonstrate that even when patients prefer chatbot encounters over human encounters, such encounters are ethically problematic. I also show that adoption of ‘empathic’ chatbots in clinical settings risks unraveling the moral fiber of clinical encounters, by robbing them of their proper existential character.

Lee, Daniel

The Promise of Moral Bioenhancement and the Implementation Challenge of Pathocracy

A B S T R A C T Several scholars have discussed Moral Bioenhancement (MBE) – biomedical interventions intended to give recipients “better” moral beliefs, dispositions, and behaviours – and how such MBE may prevent, or at least mitigate, some of the most dire existential threats that humanity and the wider biosphere of planet Earth now face. For example, Persson and Savulescu have argued consistently that MBE, either compulsorily for all or at least for most persons, is necessary in order to prevent disasters

such as nuclear war and anthropogenic climate catastrophe. This runs counter to more pessimistic voices, including Harris, who warns that MBE would not enhance morality proper, but rather would threaten our “freedom to fall”. MBE also poses issues in the broader political domain. Sparrow warns that MBE may result in a form of “moral perfectionism”, where persons are forced to conform to an externally imposed communal morality instituted by Governments. Indeed, Sparrow warns that the “morally enhanced” may be afforded more rights compared to their unenhanced peers. I extend Sparrow’s objection to compulsory MBE by articulating Andrzej Łobaczewski’s concept of “pathocracy” – a system where the psychopathic and immoral come to dominate political life, imposing their will on citizens. I discuss how the implementation of moral enhancement could be a tool of social control that psychopaths use to not only cow the general population, but also to preclude the rise of any rivals to their power. Ultimately, the implications suggest compulsory MBE is ineffective at best and extremely dangerous at worst.

LEE Ho Wai, Theo

Justifying Medical Immortality of Transhumanism

ABSTRACT Bio-enhancement enables the possibility of a radical prolonged life, or even medical immortality. However, many hold a pessimistic view as they are influenced by the Makropulos’s argument of Bernard William, that immortality is inherently bad for human beings for either the loss of identity or inevitable tedium. Yet, I argue that Bernard William misunderstood the nature of experience of human being in his argument; hence gave a mistaken conclusion. I believe immortality is instrumentally good for human beings. I argue this in two steps: (1) the correct nature of experience of human beings – temporal, allows human beings to escape the plight argued in Makropulos’ case; (2) a radical prolonged life enables human beings to flourish their life better in both quality and quantity wise. I believe the conclusion helps justify one of transhumanism’s goal of conquering aging, and offering different perspective on the effect of bio-enhancement, not just focusing on the direct boost of humans’ capability, but also the indirect effect(s) that will be brought by bio-enhancement.

Linden, Ingemar Patrick

The Growing Mismatch between Our Moral Psychology and Our Technological Power

ABSTRACT Even if it is true that, while humans sometimes benefit some individuals as an end, they never harm anyone as an end, it is here argued that it is still rational to

believe that the future for humans will be worse to the point of being no better than non-existence. For generally the benevolence of humans is severely limited to themselves and close ones like family and friends, and they care little about others, and in the pursuit of what is good for themselves and close ones, they are likely to harm others as means or as side effects. It is likely that the harms caused are greater than the benefits, since it will be seen that is practically easier to harm than to benefit, and the magnitude of harm can be much greater. The extraordinary development of the technological power of humans and the consequent explosion of the human population have magnified this asymmetry between the bad and the good to the extent that the survival of humanity is jeopardized – unless they apply their sophisticated technological power to the task of morally enhancing themselves.

Lively, Cathy

Adapting Without Capitulating: Novel Bioethical Threats and the Limits of Acceptance

ABSTRACT Bioethics is confronting novel existential threats. Artificial intelligence, emerging transplant technologies such as normothermic regional perfusion for controlled donation after circulatory death, and the rapid translation of biotechnology create new demands for professional competence. At the same time, political, religious, and ideological actors restrict access to health care, constrain individual autonomy, and target vulnerable populations. Bioethics professionals must grapple with adapting to scientific advances, while resisting the erosion of core ethical commitments. The Lancet Commission on Medicine, Nazism, and the Holocaust provides a working analytic framework: it documents how changes in laws and policies, including medical ethical codes, slowly indoctrinated healthcare professionals and scientists into an ethical reversal, not through dramatic rupture, but through incremental adoption and then normalization of the changes. This presentation uses that framework to examine current novel challenges and argues that adaptation must not mean passive acceptance. Bioethicists and healthcare professionals must distinguish between adaptation to morally permissible scientific advance and resisting policies and laws that conflict with core ethical principles. They must acquire the knowledge and skills necessary to remain competent but must also articulate and defend professional norms. In this process there is an ethical obligation to determine when routine advocacy suffices, and if at some point professional disobedience is justified. In the face of existential threats, bioethics must confront how to adapt without losing itself: how to move forward in innovation, while holding firm to the core principles of respect for persons and human rights.

Marks, Jonathan H.

War Without End: The Deadly Silence of the Law and the Responsibilities of (Bio)ethics

ABSTRACT There are laws about when countries can go to war (*ius ad bellum*) and laws about what countries can do in war (*ius in bello*). But there are no laws regarding when countries are obligated to end wars or cease hostilities. We urgently need legal norms to fill this lacuna. We also need norms to address two distinct but related issues: (a) the obligations of third parties to cease providing funding, arms, or other resources to support the armed conflicts of others, and (b) the potential for third-party exploitation of conflict mediation for their own benefit. The urgent need for such norms has been vividly demonstrated by the wars in Gaza and Ukraine. The formulation of new legal norms should be informed by these and other conflicts—both historic and continuing. But, as with public health emergency preparedness measures, these norms should be formulated to address a broad array of scenarios. Bioethicists can and should participate in the development of these norms—drawing on their expertise, access, and social status and acting collectively, as well as individually. Bioethicists' experience with principlism makes them ideally placed to contribute to the creation of norms that address these legal lacunae. Such engagement would also provide an opportunity for bioethicists to develop a broader set of principles that are truly global in origin (as well as in application) and that address quotidian as well as existential and emergent threats to individual and population health and wellbeing.

Markuckas, Marius

Reflecting on Transhumanism: Disability as an Existential Resource of Humanity

ABSTRACT The idea of transhumanism, which increasingly pervades the intellectual, cultural, and sociopolitical state of contemporary societies, elicits intense debate in both academic circles and broader social contexts concerning its true nature and ethical evaluation. On the one hand, transhumanism is considered an unprecedented opportunity for humanity to acquire new techno-scientifically grounded powers and to ensure its safe and flourishing existence in the face of various existential challenges. On the other hand, transhumanism itself is presented as a potential danger, whose threatening nature becomes particularly evident regarding certain groups of people. One such group are persons with disabilities, who, due to their impairments, seem not to fit the dominant transhumanist discourse about the necessity of creating a fully invulnerable (human) being and thereby expanding humanity's existential capacities. This presentation is aimed at reflecting on which of these two perspectives on transhumanism—either as an unprecedented opportunity for humanity or as an unprecedented danger to it—is more justified, analysing it specifically from the perspective of people with disabilities. It argues that the experience of disability can be

regarded as an existential resource that expands humanity's capacities and may therefore be successfully reconciled with certain transhumanist visions of human development.

McGowan, Michelle L. & Barbara A. Barry

Mirror, Mirror?: Disrupting the Illusion of the View from Nowhere in AI Ethics

A B S T R A C T Enthusiasm for artificial intelligence (AI) to fundamentally shift biomedical innovation is feverish. Academic medical centers are investing vast resources into AI, encouraging researchers and clinicians to embrace an ever-expanding catalog of AI tools. Nevertheless, uncritical adoption of AI raises ethical concerns. One is its potential to perpetuate what Shannon Vallor invokes is a contemporary analogue to Ovid's Narcissus, cautioning that AI's shallow portrayal of users, communities, and society can mislead and distort our reality. This distortion locates represented users in echo-chambers of reinforcing beliefs and un(der)represented users in stereotypes or invisibility. We extend Vallor's argument with Donna Haraway's critique of "the view from nowhere" to show how AI models reflect not a neutral, universal "human" vision of ourselves, but a partial, subjective view reflective of dominant groups. The guise of neutrality in AI masks power dynamics and perpetuates culturally normative hierarchies of knowledge and power. AI developers recognize this as a technical limitation – taking steps to eliminate bias from training data and models – but mitigating Vallor's Narcissus warning requires fundamentally challenging "the view from nowhere." Haraway argues all knowledge is partial and situated, and recognition of such allows us to interrogate the ethical and political dimensions of AI with appreciation for contours of knowledge production and power relations. To challenge claims of neutrality and elevate situated knowledges in biomedicine, AI developers and end users ought to amplify transparency, center varied human values and experiences, anticipate effects on end users, and promote contextually grounded AI with acknowledgement of local knowledge contributions.

Ord, Toby

Progress and Existential Risk

A B S T R A C T I show how a standard argument for advancing progress is extremely sensitive to how humanity's story eventually ends. Whether advancing progress is ultimately good or bad depends crucially on whether it also advances the end of humanity. Because we know so little about the answer to this crucial question, the case for advancing progress is undermined. I suggest we must either overcome this objection through improving our understanding of these connections between progress and human extinction or switch our focus to advancing certain kinds of progress relative to others – changing where we are going, rather than just how soon we get there.

Paquette, Erin

Rethinking Bioethical Paradigms: Centering a Systems Approach to Responsibility, Inclusivity and Justice In An Era of Rapidly Emerging Technologies and Global Existential Threats.

ABSTRACT Global existential threats—from climate destabilization and pandemics to artificial intelligence and displacement/migration—demand an ethical framework that extends beyond traditional bioethics’ clinical and research focus. This proposal argues for a paradigm shift toward an anticipatory and systems-oriented bioethics grounded in Hans Jonas’ ethics of responsibility, emphasizing our moral duty to act amid uncertainty to protect future generations. Integrating human rights, health justice, and public health ethics, this approach centers on intergenerational justice, precautionary governance, and inclusive deliberation as foundational principles for ethical decision-making across scientific, technological, and health domains. Drawing on examples including equitable vaccine distribution, governance of emerging technologies, and the ethical alignment of AI systems, the proposal demonstrates how systems thinking reveals the interconnected vulnerabilities linking ecological, economic, and social stability. By incorporating Rawlsian justice, virtue ethics, and the social contract tradition, this framework underscores that stewardship, transparency, and the promotion of truth and trust define ethical professionalism in an era of global risk. I propose that anticipatory and justice-centered bioethics be integrated into the education of innovators, engineers, scientists, and clinicians—embedding social cohesion, intergenerational flourishing, and global health equity as core professional competencies. This reframing of bioethics as an instrument of collective resilience offers a constructive and actionable response to emerging and destabilizing technologies, positioning ethics not as constraint but as catalyst for responsible innovation and human survival.

Peeler, Katherine R.

Bioethical Underpinnings in U.S. Policies Affecting the Health of Asylum-Seekers

ABSTRACT Immigrant health, including those without legal status in the United States (US), is not a new area of inquiry in ethics and bioethics. Much of the bioethics scholarship has been normative, looking at the ethics of delivering care, advocacy, and ethical analysis of the state of care in settings such as detention. This has been less so about understanding how such care came to be but rather what the ethical implications are and how future care should change to provide more ethically-sound care to immigrant populations. The aim of this paper is to describe values expressed in policies (legislation, official guidelines, and other official policy-type documents) related to the care of immigrants in U.S. federal custody over the last two decades since the inception of US

Immigration and Customs Enforcement agency. Through qualitative thematic analysis of values expressed in such policy documents, preliminary findings demonstrate that there are core values related to safety, security, and individual and public health promotion that emerge throughout a diverse sampling of policies. However, how those values are expressed and who those values are meant to apply to vary dramatically depending on the intention of the policy. Having utilized a bioethical approach of illuminating values that underpin past immigrant health-related policies, the next step in this larger project will be to apply these findings in the development of a bioethics-informed shared valued based framework to guide future policy discussions as a potentially more productive alternative to the present partisan approach to such policymaking.

Persson, Ingmar

The Growing Mismatch between Our Moral Psychology and Our Technological Power

ABSTRACT Even if it is true that, while humans sometimes benefit some individuals as an end, they never harm anyone as an end, it is here argued that it is still rational to believe that the future for humans will be worse to the point of being no better than non-existence. For generally the benevolence of humans is severely limited to themselves and close ones like family and friends, and they care little about others, and in the pursuit of what is good for themselves and close ones, they are likely to harm others as means or as side effects. It is likely that the harms caused are greater than the benefits, since it will be seen that is practically easier to harm than to benefit, and the magnitude of harm can be much greater. The extraordinary development of the technological power of humans and the consequent explosion of the human population have magnified this asymmetry between the bad and the good to the extent that the survival of humanity is jeopardized – unless they apply their sophisticated technological power to the task of morally enhancing themselves.

Rakić, Vojin

What Are Thinkers Good For in the Age of Advanced AI?

ABSTRACT There is an ancient philosophical puzzle concealed inside a very contemporary anxiety. For Aristotle, the highest form of intellect – nous poietikos – was described as separable, impassive, and eternal: thought thinking itself, free of any particular thinker. For most of the history of philosophy, this remained a speculative theological problem. Today it has become urgent. Advanced artificial intelligence systems now produce arguments, analyses, proofs, and creative works that are, by many ordinary measures, indistinguishable from the outputs of human thought. The question of whether

thought can occur without a thinker has migrated, without warning, from the seminar room into the datacenter. This talk investigates what remains irreducibly valuable about human cognition once its outputs can be reproduced without anything we would recognise as a mind. Drawing on Aristotle's account of nous and theoria, Kant's transcendental unity of apperception, and Heidegger's concept of Dasein, I argue that the human thinker cannot be reduced to a cognitive engine. Thinkers are the sites at which thought acquires stakes, is embedded in a life, is held accountable to others, and remains open to transformation through wonder, error, and mortality. I develop a concept — situated intellectual agency — to capture these dimensions, and I show that it cannot be automated, not because it is computationally complex, but because it requires a self that has something to lose. The talk then confronts the hardest objection: that if AI produces the products of thought more reliably than humans, the special features of human cognition begin to look like inefficiencies rather than virtues. I answer this by distinguishing activities valuable for their products (poiesis) from those constitutive of a flourishing life (praxis and theoria), and by examining the civilisational risks of epistemic dependence on AI systems. I conclude that the emergence of artificial thought does not signal the twilight of the thinker but rather occasions a deeper understanding of what thinking is — and that navigating this moment well is itself one of the most pressing intellectual tasks of our time.

Ravitsky, Vardit et al.

What are Humans Good for in the Age of AI?

A B S T R A C T AI-powered chatbots can at least mimic and perhaps fully perform relationship roles that, until recently, have been the sole domain of human beings. They raise the question of whether there is any *in-principle* reason to constrain the use of chatbots in relationship roles. In this paper we trace the strengths and weaknesses of several possible claims about what humans bring to relationships that is morally special and might support an in-principle constraint. *Species membership* is largely unsatisfactory, we argue. *Embodiment* and *personhood* are more promising but still surmountable. We propose that *shared moral governance* is an activity in some relationship roles that should be regarded as morally special and protected from AI. We then consider how implications for the use of chatbots would be worked out.

Recasens Subias, Josep

Conditional Protection for Climate Disasters

A B S T R A C T This paper asks whether compensation for individuals who suffer

climate-related disasters should be conditional on their mitigative and adaptive actions. Most philosophical accounts of Loss and Damage (corrective and distributive justice) explain why there is a claim to compensation but say little about its limits. Unconditional aid can also entrench vulnerability by enabling rebuilding in harm-prone areas. I develop a responsibility-sensitive framework that integrates sufficientarian aims with Dworkin's hypothetical insurance, still not applied to disaster response. The framework distinguishes backward-looking conditionality (responsibility for avoidable risk-taking) from forward-looking conditionality (reasonable requirements attached to the use of aid), and evaluates both against tests of fairness, feasibility, and incentive effects. The view of Conditional Protection holds that (i) everyone is owed protection up to a sufficiency floor; (ii) deductions based on responsibility apply only when agents had genuine, affordable alternatives; (iii) forward-looking conditions must be safety-enhancing (e.g., safer siting, energy-efficient reconstruction), proportionate, and paired with enabling support (e.g., relocation subsidies); and (iv) structural disadvantage triggers exemptions and priority. I illustrate the approach with cases of coastal housing and drought-affected pastoralists, and derive design implications for cash transfers, relocation packages, and index-based insurance (including risk-pool participation and deductibles calibrated to opportunity). The account avoids moral hazard without blaming the blameless, reconciles compensation with mitigation and adaptation priorities, and offers action-guidance for Loss and Damage policy.

Rekers, Romina

Moral preparedness for health emergencies through transdisciplinary reflective equilibrium

ABSTRACT Open moral knowledge now, save later!

Public health emergencies create unprecedented moral complexity that demands the continual revision of our moral commitments. As health systems increasingly institutionalize preparedness as a permanent function, moral preparedness remains underdeveloped. This presentation argues for recognising moral preparedness as a necessary dimension of health emergency preparedness. It introduces Transdisciplinary Reflective Equilibrium (TRE) method as a structured approach to revising moral commitments and enhancing moral preparedness in public health emergencies through the co-production of moral knowledge with practitioners and the public. Developed through collaborations between the University of Oxford and Johns Hopkins University (GLIDE), and currently expanded through WHO-funded research, TRE integrates philosophical analysis with participatory deliberation. The method comprises four key elements: 1. structured procedure: an ethics lab bringing together diverse disciplinary and stakeholder perspectives, including non-academic participants; 2. Normative outcome: the co-production of moral knowledge, rather than the mere description of moral

preferences or reasoning patterns; 3. Causal and constitutive benefits: both the deliberative process and its outcomes strengthen participants and their communities. For example, non-academic participants gain the ability to construct more justified moral arguments simply through participation; 4. Equal control: academic and non-academic participants share equal control over the research process, including the right to introduce new cases and to contest judgments or proposed revisions of principles. The presentation discusses how TRE can contribute to moral preparedness, thereby strengthening the moral and institutional resilience of public health systems for facing emergencies.

Resch, Marie-Kristin

Moral Enhancement – A Remedy for Democratic Regression?

A B S T R A C T Today's liberal democracies are facing a profound crisis. Rising polarization and growing anti-democratic tendencies are eroding mutual recognition and cooperation as the very foundations of democratic life. Against this backdrop, the idea of moral enhancement – the biomedical improvement of moral capacities such as empathy, altruism and a sense of justice – appears to be a potential response to democratic regressions. Strengthening our moral capacities could reduce hostility between different social groups, restore mutual recognition, and thus counteract polarization and division. Yet, this proposal also raises far-reaching ethical and political questions: How can we ensure that moral enhancement truly achieves its intended effects, extending our enhanced moral capacities to other social groups rather than just being limited to our own ingroup, thereby potentially reinforcing existing group-based divisions? And how can we implement moral enhancement in a democratic manner in order to provide a justifiable solution to the diagnosed democratic regressions? Examining these questions allows us to rethink the relationship between biotechnology, morality and democracy, and to identify whether moral enhancement could serve as a remedy for the crisis facing liberal democracies – or whether it risks deepening the very challenges it aims to overcome.

Ruiz Serrano, Mario

Biopolitical Prelude to the End of Conscious Alienation: Tortuous Rhetorical Choreographies in the Recommendations of an International Commission (NAM, NAS & RS, 2020) on the Instrumentalization of Humans Mediated by Cas9–gRNA Ribonucleoproteins

A B S T R A C T In 2020, the International Commission on the Clinical Use of Human Germline Genome Editing (NAM, NAS & RS) proposed the first institutional framework

that considers the creation of genetically modified humans ethically and socially acceptable under certain conditions. This work critically analyzes the recommendations of that commission, questioning the logical and empirical soundness of its arguments and the biopolitical framework that sustains them. Through a historical, philosophical, and bioethical methodology —based on the principles of Beauchamp and Childress, and Carole Pateman’s democratic theory— the notions of informed consent, democratic legitimacy, and sufficiency of preclinical evidence are examined. The analysis shows that the fundamental assumptions of the report are inconsistent: the extrapolation from animal models to humans lacks a solid basis (there are physical and moral constraints around clinical testing with human embryos); parental consent does not substitute for the consent of the nonexistent subject; and the democratic mechanisms invoked do not guarantee effective representation or citizen control. The study concludes that the report’s recommendations are logically invalid and empirically unfounded, and that their adoption would imply the normalization of a potentially authoritarian technocratic governance over human biology. The work proposes, in contrast, an ethics of radical democracy and a reevaluation of collective political judgment as a prior condition for any clinical application of germline genome modification technologies.

Sandberg, Anders

Hinges of History and of Technology: When Can We Know Enough to Act on Existential Risk?

A B S T R A C T Differential technology development (DTD) — preferentially accelerating protective over dangerous technologies — is a key strategy for existential risk mitigation. But DTD has an underappreciated epistemic precondition: identifying, prospectively, the critical periods when intervention is both consequential and feasible. I call these technological hinges. I propose a three-condition framework — consequentiality, malleability, and legibility — and examine cases from nuclear weapons to AI where these conditions were variously satisfied or missed. The analysis suggests that the binding constraint on managing existential risks is not normative but epistemic: civilisations need not just better principles but better instruments for knowing when to act.

Savulescu, Julian

Ethics of Polygenic Editing of Embryos

A B S T R A C T In November, 2018, Chinese scientist He Jiankui announced the birth of the world’s first gene edited babies, Lulu and Nana. I will review the circumstances

around this remarkable announcement. I will discuss the ethics of polygenic editing of human beings and compare this to genetic selection of embryos. I will argue that there is a moral obligation to pursue research into human gene editing and a moral obligation to select the best child (procreative beneficence). I will outline a translational pathway for ethically justified first in human trials of gene editing. I will evaluate recent calls for a moratorium on gene editing of human beings and compare gene editing to genetic selection. Finally, I will argue that the birth of Lulu and Nana heralds the arrival of an enhancement revolution. I will discuss how this revolution could proceed ethically. I will discuss briefly the prospect of human moral bioenhancement.

Shevchenko, Sergei

The Double Edge of Biological Multirealizability: Limiting AI Control, Enhancing AI Risk Generation

A B S T R A C T Multirealizability – the capacity of biological systems to realize the same function through multiple structural configurations – is a central concept in theoretical biology (Koskinen 2019). It contributes to the resilience of living systems while simultaneously complicating efforts to regulate or control them. In this respect, AI-assisted design of therapeutic agents encounter multirealizability as a fundamental constraint at the input stage. The present inability to account for all possible structural realizations of a to-be-controlled pathological process constitutes a major limitation on the biomedical AI use. Meanwhile, at the output stage, AI systems are capable of generating novel structural realizations of existing functions – for instance, rewriting a text while preserving its meaning, or designing protein sequences with different amino acid compositions but similar biological activity. This generative capacity promises more effective therapeutic design, but also entails new existential risks. One such risk involves AI-enabled creation of biohazards, whose genetic or amino acid sequences are absent from existing databases and therefore cannot be automatically restricted (Wittmann et al. 2025). Proposed safeguards include the establishment of a global tracking system for all synthetic biology orders (Bloomfield et al. 2024), though such coordination appears implausible in the context of growing international polarization. A broader solution would aim to align AI’s monitoring capacities at the input stage (recognizing function from structure) with its generative capacities at the output stage (creating multiple structural realizations of a function). This alignment could mitigate AI-related biological risks, through deepening the connection between bioethics and theoretical biology.

Szocik, Konrad

Survivalism and Justice: A Feminist Bioethical Assessment of Space Enhancement and Reproduction under Existential Risk

A B S T R A C T Contemporary proposals for off-planet settlement and biomedical enhancement are frequently justified by survivalist and longtermist arguments. This paper evaluates their moral status from the perspective of feminist bioethics. First, we analyze spaceflight as a structurally coercive context in which consent and autonomy are relationally constrained; mission infrastructures, selection procedures, and closed hierarchies distribute risks and benefits asymmetrically. On this basis, we assess moral bioenhancement and germline editing for long-duration missions, highlighting uncertainties about behavioral effects, irreversibility, and the governance problems created by mixed modified/unmodified populations. We then examine whether space missions provide special justification for enhancement given isolation, small population size, and collective-action dynamics. While these features may lower coordination barriers, they simultaneously heighten risks of domination, positional disadvantage, and consent deficits. Finally, we locate space projects within terminal ethical scenarios such as reconciliation, escape, and reproductive restraint and argue that permissibility turns not on feasibility or speculative legacy, but on satisfaction of substantive justice conditions: non-domination, reparative and distributive justice, and reproductive justice. Where these conditions cannot be met, reproductive non-continuation and non-expansion are normatively defensible responses that prioritize the avoidance of imposed harm over preservation at any cost. The contribution is a justice-first framework for evaluating enhancement and reproduction in existential-risk governance: continuity has no lexical priority over protecting present and foreseeable persons in demonstrably constrained environments.

Tang Ying Fei

The New Fashion in Humanity

A B S T R A C T Psychedelics to treat mental disorders, surrogate robots to help women bear their children, digital afterlife to manage grief ... Are these the essence of human wisdom and kindness to improve lives, or more of human ambition and rebellion to break the rules and achieve the impossible? When Artificial Intelligence extends far beyond school or work assistance and daily problem-solving to advanced surgical robots connecting in real-time across countries, to potentially influencing life-and-death matters, is it a troublemaker or a promising work in progress deserving our investment? What kind of world awaits us? What roles and responsibilities do the younger generations hold in the creation, execution, regulation, and mastery of these technology monsters? In my humble opinion, the solution is really to recognize the limitations of both human and

AI, to stay open-minded, to advance and polish but not to forget the core of humanity. Reflecting on history – “Man a Machine” theory was raised by Julien Offray de La Mettrie in 1747; Jacques Offenbach's opera “Les contes d'Hoffmann”, which premiered in 1881, narrates the story of a protagonist who cannot tell if the woman he loves is a real opera singer or a mechanical doll; Tilly Norwood, the first AI-generated actress was born in 2025. Perhaps the iconoclastic, avant-garde fashion in enhancement is itself a manifestation of what is rooted in our human nature – the desire to go beyond limits. In face of challenges, criticisms and confusion, can human wisdom transform chaos into a new rhythm in harmony?

Vinogradova, Olga

Challenges Experienced by Swiss Healthcare Professionals during COVID-19: A Qualitative Study

ABSTRACT The qualitative study “Decision Making in Times of Scarce Resources: A Mixed-Method Study” investigates healthcare workers’ perceptions, experiences, and challenges in implementing restrictive measures during the COVID-19 pandemic in Switzerland, with a special focus on protecting older adults. This research is conducted within the broader context of resource limitations, preferences for resource allocation, and care satisfaction among older adults throughout the pandemic. To carry out the study, fifty-seven semi-structured interviews were conducted with healthcare workers across five Swiss cantons. The interview data were analyzed thematically to identify core issues related to restriction measures. Participants were selected according to two main criteria: employment in the healthcare sector during the COVID-19 pandemic in Switzerland, and involvement either in developing guidelines or policies for older adult care, or applying these guidelines in practice. These criteria reflect the aims of the larger research project, which seeks to understand how key stakeholders managed the care of elderly patients and pre-hospital triage under conditions of resource scarcity. In this presentation, the various challenges faced by healthcare providers and stakeholders in pre-hospital and in-hospital triage in resource-limited settings will be discussed. These include institutional and organisational challenges such as communication issues, changes in patient care and workloads, reorganization tasks, team management, and balancing overstaffing and understaffing. The findings contribute to a deeper understanding of the implications of triage decisions during public health emergencies.

Walker, Alexis

Existential Threat Discourses, Collective Organizing, and University Futures

ABSTRACT A considerable literature has emerged theorizing universities as under existential threat worldwide—from authoritarian incursions on academic freedom to austerity politics and market logics that erode public missions, and from climate change and technological developments that challenge the very conditions of collective existence. Yet the framing of “existential threat” itself does political and analytic work. As Daniel Goldstein and others have argued in contexts ranging from security to finance, discourses of threat and crisis not only describe external dangers but also reorder priorities, distribute vulnerabilities, and shape modes of authority. Drawing on interviews and document analysis of graduate and contingent faculty organizing within U.S. universities, this paper situates these debates within a broader concern with liberal institutionalism. It shows how existential threat discourses can channel energies into preserving forms already embedded with inequality, where exclusion and precarity have been structured into everyday governance and labor relations. At the same time, collective organizing emerges as a means through which endurance takes political form. Reading these materials through anthropological theories of endurance under late-liberal abandonment highlights not only the informal practices and solidarities that sustain collective life amid dispossession, but also the affective and political labors that maintain the possibility of renewal and resistance. Here, endurance is an ambivalent process—one in which institutions may persist through retrenchment, yet internal dissent can endure as well, cultivating the conditions through which new futures might take hold.

Walker, Rebecca

One health one ethic? Unified research oversight for troubled times

ABSTRACT The COVID-19 pandemic and global climate crisis each underscore how health is interconnected across species. The One Health movement aims to leverage these sites of mutual impact for shared benefit. Such benefit depends on testing and developing health interventions using both human and nonhuman animal subjects. Yet, as some public health advocates and practitioners embrace a unified view of health in the face of global threats, the ethical systems of research oversight have lagged. Worldwide, human and nonhuman animal studies are managed through divergent oversight systems functioning under distinct ethical principles. While typically unquestioned, employing disunified ethical systems to guide research is conceptually remarkable. Given unprecedented existential threats calling for novel science solutions, our research oversight systems must stand up to ethical scrutiny. This talk critically assesses key potential justifications for the current divergence in research ethics oversight. Common rationales for treating human and nonhuman subjects of science differently have included appeals to moral status and welfare differences. I focus instead on a possible rationale stemming from structural differences in the science of human and nonhuman animal

studies. While this rationale is more promising than those offered by moral status or welfare, I argue it ultimately fails because of cross-over in human and animal studies. Instead, we require a unified ethical system that can both do justice to the complex overlap of human and nonhuman animal lives and capacities and allow for significant structural diversity in the science of biomedical research.

Zhong, Weian

Euthanasia, Absolute Duties, and Moral Permission

A B S T R A C T Debates about euthanasia often miss their target. Proponents offer reasons why euthanasia can be morally permissible, while opponents insist it should never be permissible; no amount of added reasons or safeguards can satisfy a categorical prohibition. This article addresses the stalemate by rejecting the assumed existence of absolute duties against involvement in euthanasia. I argue that such duties fail at all three perspectives typically invoked: (1) patients do not bear an absolute duty never to commit suicide; (2) medical practitioners do not bear an absolute duty never to perform euthanasia; and (3) the state does not bear an absolute duty never to legalize it. At the patient level, choices for death are not invariably irrational or involuntary, “clear windows” allow capacity assessment, and life’s intrinsic value does not entail final or absolute value. At the practitioner level, risks to trust and error are real but mitigable; excluding clinicians can increase overall harm by pushing patients to unsafe, unregulated means; and medicine’s life-preserving ethos coexists with a duty to relieve suffering, with conscientious refusal protected. At the state level, slippery-slope and discrimination worries fade when legalization is tied to multiple conditions—capacity, voluntariness, informed and valid consent, clinical assessments, proportionality, independent review—alongside robust reporting, equitable palliative care, and safeguards for the vulnerable. The upshot is conditional pluralism: a framework that honors life’s worth and palliative care’s centrality while allowing, in circumscribed cases, a competent patient’s considered judgment to prevail. Euthanasia remains morally significant and calls for humility, not absolutism; but the case for a categorical ban fails.

Zuk, Peter

Mental Privacy, Self-Expression, and Hermeneutical Injustice

A B S T R A C T I begin by examining the two major approaches to mental privacy, *control theories* and *contextual integrity*. Neither of these, I argue, supplies an account of what is distinctive of mental privacy and of neurotechnological modes of apprehending mental states. I propose that neurotechnological “mind-reading” is distinctive in entirely bypassing verbal and behavioral expression of mental states. My explanation of why this

makes a *normative* difference relates the value of mental privacy to that of self-knowledge and self-expression. In externalizing mental states in a way that bypasses person-level inputs, neurotechnological mind-reading involves a form of automatic (and, if compulsory, *forced*) expression that competes with our natural, analogue capabilities for unmediated self-expression. Neurotechnological mind-reading also threatens our self-knowledge when it offloads introspective awareness of our own mental states to an external device. Relying on neurotechnological “exospection” of this sort to know our own minds risks eroding our analogue capability to know them through good-old-fashioned introspection, and thereby inflicting what Miranda Fricker calls *hermeneutical injustice*. Reliance on an external interpretation of one’s inner mental life threatens to impoverish the conceptual repertoire through which we understand our own experience. And if we understand subjectivity as constituted not only by first-personal phenomenal experience, but also in part by first-personal interpretation of that experience, neurotechnological mind-reading risks diminishing the richness of our subjectivity itself. This potential diminishment of subjectivity implicates our humanity itself, and thereby rises to the level of a genuinely *existential* risk. Rather than mitigating existential risk, neurotechnology may itself present such a risk.