FUNDING

This work is supported through the Bridge2AI program, NIH Common Fund, grant numbers 1OT2OD032720-01 and 1OT2OD032742-01.

ORCID

Jean-Christophe Bélisle-Pipon http://orcid.org/0000-0002-8965-8153
Vardit Ravitsky http://orcid.org/0000-0002-7080-8801

REFERENCES

Barnes, C., M. R. Aboy, T. Minssen, J. W. Allen, B. D. Earp, J. Savulescu, and S. P. Mann. 2025. Enabling demonstrated consent for biobanking with blockchain and generative AI. *The American Journal of Bioethics* 25 (4):96–111. doi: 10.1080/15265161.2024.2416117.

Bélisle-Pipon, J. C. 2024. Why we need to be careful with LLMs in medicine. *Frontiers in Medicine* 11:1495582. doi: 10.3389/fmed.2024.1495582.

Bélisle-Pipon, J. C., E. Monteferrante, M. C. Roy, and V. Couture. 2023. Artificial intelligence ethics has a black box problem. *AI & SOCIETY* 38 (4):1507–22. doi: 10.1007/s00146-021-01380-0.

Bélisle-Pipon, J. C., M. Powell, R. English, M. F. Malo, V. Ravitsky, Y. Bensoussan, and Bridge2AI–Voice Consortium. 2024. Stakeholder perspectives on ethical and trustworthy voice AI in health care. *Digital Health* 10. doi: 10.1177/20552076241260407.

Evans, B. J., and A. Bihorac. 2024. Co-creating consent for data use—AI-powered ethics for biomedical AI. *NEJM AI* 1 (7):AIpc2400237. doi: 10.1056/AIpc2400237.

Rhim, J., H. Gallois, V. Ravitsky, and J. C. Bélisle-Pipon. 2024. Beyond consent: The MAMLS in the room. *The American Journal of Bioethics* 24 (10):85–8. doi: 10.1080/15265161.2024.2388737.

THE AMERICAN JOURNAL OF BIOETHICS
2025, VOL. 25, NO. 4, 115–118
https://doi.org/10.1080/15265161.2025.2470694

Taylor & Francis
Taylor & Francis Group

OPEN PEER COMMENTARIES

OPEN ACCESS Check for updates

On the Complexities of Enabling Demonstrated Consent

Panagiotis Alexiou, Joel Azzopardi, Claude Julien Bajada, Jean-Paul Ebejer, Gillian M. Martin and Nikolai Paul Pace

University of Malta

Barnes et al. (2025) introduce a novel vision for biobanking consent in their article "Enabling Demonstrated Consent for Biobanking with Blockchain and Generative AI." Their concept of "demonstrated consent" leverages blockchain technology and non fungible tokens, along with large language models to enhance transparency and participant engagement. We put this novel framework in context of existing approaches, and highlight some key questions that arise.

The central question that Barnes et al. seek to address is the inherent limitations of the following two traditional consent models in biobanking, namely:

- **Study-specific consent:** A type of consent where individuals give permission for the use of their bio-samples and data to be used in a single, well defined, research project. While ethically robust, it is often impractical in the context of long-term biobanking due to the increased administrative burden and inflexibility.
- **Broad consent:** A type of consent where individuals allow their bio-samples and data to be used in future, sometimes unspecified, research projects, with few or no specific restrictions. Importantly, without the need to be re-contacted

CONTACT Panagiotis Alexiou panagiotis.alexiou@um.edu.mt Centre for Molecular Medicine and Biobanking, University of Malta, Msida MSD 2080, Malta.

or consulted. This approach is more efficient, but can undermine the autonomy of participants by failing to provide sufficient information about future research uses. Broad consent needs to be paired with strategies of risk mitigation, and continuous provision of information to participants.

In response to these challenges, various alternative models have been proposed, including tiered informed consent (Tiffin 2018), meta-consent (Ploug and Holm 2016), and dynamic consent (Kaye et al. 2015; Budin-Ljøsne et al. 2017) These models try to increase the involvement of participants, but are vulnerable to issues similar to study-specific consent. Dynamic consent in particular, has gained traction as a means of allowing participants to dynamically update their consent preferences in real time, thus tailoring their participation to studies based on their preferences. Individuals however, have to constantly manage their consent, leading to potential choice overload and consent fatigue. The "demonstrated consent" model proposed by Barnes et al. aims to circumvent these issues by providing a secure, transparent, and easily accessible source of information, without requiring participants to continuously manage their preferences. The central contribution of this manuscript is the proposed integration of non-fungible tokens (NFTs) and large language models (LLMs) to tackle these issues.

Blockchain technology provides immutability, transparency, and auditability of data, creating a permanent record of sample use. NFTs are unique digital representations of each donated sample, and their metadata can contain information about the sample's characteristics, storage location, and associated consent preferences. However, an issue with blockchain technology involves the "right to erasure" under the GDPR (Mamo et al. 2020), due to the blockchains immutable structure. Zero-knowledge proofs and homomorphic encryption are proposed as potential future solutions to such concerns, but without a specific implementation proposed. The practical implementation significantly influences key aspects of the proposed system, such as participant roles in updating the blockchain and the choice of consensus mechanism (e.g. proof-of-work, proof-of-authority, or proof-of-stake).

The second innovative point of Barnes et al. is the proposed use of LLMs to enhance the accessibility and comprehension of complex information related to sample usage. These customized LLMs would be able to present information in an understandable, interactive way. A "chat with your samples" application that would enable lay-person participants to engage with complex data such as study protocols, ethical provenance documents, and so on. LLMs can indeed reduce the explanatory burden on participants, but can be prone to hallucination (Survey of Hallucination in Natural Language Generation 2023) which could make their deployment difficult at this point in time. The use of retrieval-augmented generation (RAG) can help the LLM produce more grounded and factual data, however the risk is not fully eliminated. In such a sensitive setting, it is unclear what the acceptable error/hallucination rate is, and where the ethical balance between producing understandable engagement and absolutely factual information lies. Given the large amounts of documentation, any intermediary interpretation of the documentation is liable to produce factual errors, a fact that applies also to human interpretation but can be exacerbated by current LLMs' tendency to hallucination.

Deploying a LLM requires either self hosting which comes with certain hardware requirements, or else license requirements to use a commercially available LLM via API. Local hosting involves a tradeoff between cost and effectiveness of the LLM, while using commercial LLMs via API creates another point of potential privacy and security failure. Care must be taken that such a solution does not exclude smaller biobanks in resource-limited settings, but also does not create new privacy issues by sharing information with LLM providers. The use of an LLM for obtaining consent is not without its challenges. In dynamic consent, participants' understanding is often assessed through a brief quiz that demonstrates their grasp of the study and their rights/preferences. How, then, would an LLM evaluate their comprehension of what their consent entails?

Another concern is that the proposed interaction between participants and the blockchain being facilitated via an LLM in effect hides the blockchain behind an LLM interface, undermining the auditability that the blockchain is meant to provide. The core purpose of the blockchain is to allow an immutable audit trail of participants' consent changes. If the end user cannot directly access and verify the data stored on the blockchain, then this advantage is lost. The LLM middle layer, while user-friendly, creates a barrier between the donor and the blockchain. Additionally, the proposed private, permissioned blockchain operated by a service provider, reinforces the idea that the participant would not have direct access to the blockchain. In this case, why does the data need to be stored in a blockchain, and not a simple relational database, if the blockchain will not be exposed to the participant? Could a hybrid model be implemented, in which the LLM is the user-friendly primary mode of access for participants, but the

blockchain is still accessible via a secondary interface for auditing and validation reasons?

Further, we would like to consider the question of whether LLM/NFT mediated transparency would actually foster trust in participants. It would be an interesting question to look at communities that are classically under-represented in large-scale biobanking initiatives, communities that have different levels of background genomic literacy, as well as those that have a historical mistrust of medical research (Angelo et al. 2022). The "demonstrated consent" model is designed to foster trust by providing participants with a tangible sense of their contribution to the scientific process. However, complex systems such as the one proposed here could be perceived as opaque and overly complicated, potentially undermining rather than enhancing trust, especially in communities where a historical mistrust of medical research is present. Additionally, whilst the use of LLMs is being considered on the premise that they help induce trust, there is also the possibility that communities may actually distrust AI systems. Therefore, there is a need for testing such assumptions on the effectiveness of such systems and how well they are understood and accepted by potential donors, especially in diverse types of communities.

The authors additionally introduce a "balance criterion" for the evaluation of ethical consent frameworks, beyond the three previously proposed (Mikkelsen et al. 2019), namely: information, value, and duration. The concept of the "balance criterion" is meant to balance the tension between individual autonomy and collective good. While study-specific consent prioritizes individual autonomy to the detriment of the research progress, conversely broad consent enables more efficient research but may not as fully respect individual autonomy. The proposed "balance criterion" seeks to find a middle ground between these two, while also incorporating the concept of justice. It was not clear from the original paper, how much direct involvement of the participants with the LLM is needed in order to affect their consent status. Assuming that broad consent is the baseline, would a participant who does not interact with the LLM remain in that status, or would they be required to engage and review periodically?

In conclusion, the "demonstrated consent" model proposed by Barnes et al. offers a novel approach to biobanking consent, using blockchain, NFTs, and LLMs to enhance transparency and participant engagement. The use of blockchain technology provides immutability and auditability, but certain technical challenges remain, such as ensuring compliance with "right to erasure." Accessing the underlying blockchain via an LLM layer can allow participants better comprehension of complex information and hands on interaction with the use history of their data, but could also obfuscate or mislead the users, at least until the risk of LLM hallucinations is completely eliminated. Hence, it is suggested that the LLM layer is presented as an optional functionality whereby users can opt-out of its use. The proposal of keeping the blockchain private also defeats the purpose of using blockchain technology in the first place, and may need to be rethought into a hybrid model where the blockchain is exposed, but an additional LLM based system can be used to increase participant engagement. Finally, the introduced "balance criterion" is an interesting idea, and it would be interesting to see further research on how "demonstrated consent" actually enhances this criterion. In summary, the "demonstrated consent" model offers a promising direction for biobanking consent, but requires careful consideration of both its technical and ethical implications. The central challenge will be to ensure that these systems are both transparent and trustworthy, and that they work for all types of communities while also supporting the wider societal aims of biobanking.

DISCLOSURE STATEMENT

No potential conflict of interest was reported by the author(s).

FUNDING

This work was supported by the University of Malta's Data Integrity and Stewardship Cluster (DISC), and the project BioGeMT (HORIZON-WIDERA-2022 Grant ID:101086768) at the University of Malta.

ORCID

Panagiotis Alexiou http://orcid.org/0000-0003-3437-7482
Claude Julien Bajada http://orcid.org/0000-0001-6138-4851
Gillian M. Martin http://orcid.org/0000-0002-5281-8117

REFERENCES

Angelo, F., D. Veenstra, S. Knerr, and B. Devine. 2022. Prevalence and prediction of medical distrust in a diverse medical genomic research sample. *Genetics in Medicine* 24 (7):1459–67. doi:10.1016/j.gim.2022.03.007.

Barnes, C., M. R. Aboy, T. Minssen, J. W. Allen, B. D. Earp, J. Savulescu, and S. P. Mann. 2025. Enabling demonstrated consent for biobanking with blockchain

and generative AI. *The American Journal of Bioethics* 25 (4):96–111. doi: 10.1080/15265161.2024.2416117.

Budin-Ljøsne, I., H. J. A. Teare, J. Kaye, S. Beck, H. B. Bentzen, L. Caenazzo, C. Collett, F. D'Abramo, H. Felzmann, T. Finlay, et al. 2017. Dynamic Consent: A potential solution to some of the challenges of modern biomedical research. *BMC Medical Ethics* 18 (1):4. doi:10.1186/s12910-016-0162-9.

Kaye, J., E. A. Whitley, D. Lund, M. Morrison, H. Teare, and K. Melham. 2015. Dynamic consent: A patient interface for twenty-first century research networks. *European Journal of Human Genetics* 23 (2):141–6. doi:10.1038/ejhg.2014.71.

Mamo, N., G. M. Martin, M. Desira, B. Ellul, and J.-P. Ebejer. 2020. Dwarna: A blockchain solution for dynamic consent in biobanking. *European Journal of Human Genetics* 28 (5):609–26. doi:10.1038/s41431-019-0560-9.

Mikkelsen, R. B., M. Gjerris, G. Waldemar, and P. Sandøe. 2019. Broad consent for biobanks is best – provided it is also deep. *BMC Medical Ethics* 20 (1):71. doi:10.1186/s12910-019-0414-6.

Ploug, T., and S. Holm. 2016. Meta consent – A flexible solution to the problem of secondary use of health data. *Bioethics* 30 (9):721–32. doi:10.1111/bioe.12286.

Survey of Hallucination in Natural Language Generation 2023. *ACM Computing Surveys*. New York, NY: ACMPUB27. doi: 10.1145/3571730.

Tiffin, N. 2018. Tiered informed consent: Respecting autonomy, agency and individuality in Africa. *BMJ Global Health* 3 (6):e001249. doi:10.1136/bmjgh-2018-001249.

THE AMERICAN JOURNAL OF BIOETHICS
2025, VOL. 25, NO. 4, 118–120
https://doi.org/10.1080/15265161.2025.2470683

Taylor & Francis
Taylor & Francis Group

OPEN PEER COMMENTARIES

OPEN ACCESS Check for updates

Consent, Legal Certainty and the Need for Governance

Benjamin Bartlett, Catherine Bowden, Sarah Devaney and Søren Holm

University of Manchester

In their paper Barnes and coauthors describe a consent model Demonstrated Consent for biobanking based on a blockchain structure with ongoing communications and consent requests with participants generated through the use of generative AI, e.g. a large language model (LLM) (Barnes et al. 2025). When used in relation to a single specific biobank, the proposed model is essentially an extension of Dynamic Consent, providing more personalized interaction with research participants than current Dynamic Consent models do. The proposed model is, however, intended to function for an "inter biobank network" at "any level of decentralization" (Barnes et al. 2025, 103).

In the following we will briefly discuss two ethico legal issues that are raised by the proposed model, based on our own research on blockchain enabled consent for health data sharing and on data trusts (Bartlett et al. 2024, Cunningham et al. 2021, 2022, Neumann et al. 2023, Redhead et al. 2025).

The first issue is created by the tension between personalization, precision and uniformity in relation to AI generated consent requests and the consent subsequently provided by participants. The authors write:

> Importantly, LLM integration could enable the *personalized* provision of information. Just as LLMs can be adapted to generate text in a specific style, so can they be adapted to provide text in a format that is distinctively engaging to, and understandable by, each individual prospective donor. For instance, this could involve the adaptation of complex consent information into simpler language or the use of analogies, case scenarios, and metaphors to convey the nuances of biobanking research and its potential implications. Such a personalized approach not only demystifies the consent process but also empowers donors to make informed decisions that align with their values and expectations. (Barnes et al., 105, emphasis in original, references removed)

The problem here is that whereas the instructions to the LLM may be precise in relation to exactly what the consent covers, the personalization process will entail that each participant is presented with different

CONTACT Søren Holm soren.holm@manchester.ac.uk Department of Law, Centre for Social Ethics and Policy, University of Manchester, Manchester M13 9PL, UK.