
Plan Overview

A Data Management Plan created using DMPonline

Title: Automating the Hard Choices: Agentic AI's impact on Personhood and Moral Identity

Creator: Jessica Ciupa

Principal Investigator: Vaishak Belle

Affiliation: University of Edinburgh

Template: UoE Default DMP template for PGRs

ORCID ID: 0000-0001-5573-8465

Project abstract:

As agentic AI systems increasingly occupy safety-critical domains, discourse has focused on trustworthiness and interpretability. Yet a deeper concern remains underexplored: the potential erosion of personhood through the automation of difficult, morally weighted decisions. This paper argues that free will expressed through hard choices, particularly ethical dilemmas, is foundational to self-formation and human flourishing. Drawing on Kierkegaard's conception of choice as self-actualisation, we contend that removing existential weight from human decision-making may undermine the very processes through which identity and moral character are formed. We distinguish this claim from the existing deskilling and automation complacency literature: what is at risk is not only moral competence but also moral identity. To empirically explore this inquiry, we present a preliminary qualitative study examining the attenuation of the impulse to intervene as an early warning sign of identity-level harm. Using questionnaires and semi-structured interviews, we explore whether repeated delegation of difficult decisions to autonomous agents progressively extinguishes the act of intervention, but also the impulse to intervene at all. We analyse shifts in moral ownership, paying particular attention to accounts in which the impulse to self-author a decision has faded or become difficult to recall. We inquire whether this is a measurable precursor to epistemic surrender and moral disengagement, and ultimately to deeper erosion of moral identity. This paper reflects on the technical contributions of AI to decision-making in moral dilemmas, considering the potential consequences through the philosophy of mind, moral psychology, and human-AI interaction.

ID: 202437

Start date: 01-05-2026

End date: 24-05-2026

Last modified: 17-04-2026

Grant number / URL: <https://www.cdt-d2air.uk/>

Copyright information:

The above plan creator(s) have agreed that others may use as much of the text of this plan as they would like in their own plans, and customise it as necessary. You do not need to credit the creator(s) as the source of the language used, but using any of the plan's text does not imply that the creator(s) endorse, or have any relationship to, your project or proposal

Automating the Hard Choices: Agentic AI’s impact on Personhood and Moral Identity

Administrative Information

1) School or Institute

- CSE - School of Informatics
- School of Informatics funded by the centre of doctoral training programme for dependable and deployable AI for robotics (CDT - D2AIR), in the institute of Applications Institute of Artificial Intelligence (AIAI).

2) Name and Contact details of supervisor(s)

Dr. Vaishak Belle
vbelle@ed.ac.uk

3) Project start date

2026-05-01

4) Project end date

2026-05-24

Data Collection

5) Data Collection

Data will be collected majority on JISC surveys. However, the simulation will be an external link with only the metric of interruption of the moral decision recorded.

Mixed-methods design and both primary quantitative and primary qualitative data. Quantitative behavioural timing measurement (intervention latency), and following likert-scale response for reasonings for interruption. Preliminary questionnaire for capturing habitual AI tool use.

Data will be logged and recorded in CSV format prior to analysis. Processed Data will be provided and stored on DataShare (Behavioural Timing Data, Likert Scale, Qualitative Reflection, Analysis scripts saved in the repository). README file for the data repo will be used, with versioning of all transformation and updates recorded. Participants will be kept anonymous. This study does not rely on data sourced from third parties with no access to restricted datasets, repositories or externally held records. No data sharing agreements, access contracts or use licenses are required at this collection stage.

Documentation & Metadata

6) Documentation & Metadata

A README.txt will be used in the data repository with descriptions of the raw, processed, qualitative data and the scripts required to run them.

Ethics & Legal Compliance

7) Ethics & Legal Compliance

This study will be submitted for ethical review to the relevant School Research Ethics Committee prior to data collection commencing. As the study involves human participants, all participants will be provided with a participant information sheet and will be required to sign an informed consent form before taking part; completed consent forms will be stored securely alongside the collected data. Although this study does not collect sensitive or personally identifiable data, participants are identified by anonymous IDs only, all data collection, processing, and storage procedures have been designed in compliance with the Data Protection Act 2018 (DPA2018/GDPR) as it applies to research. The lead researcher (s1810129@ed.ac.uk) will complete the University's mandatory Data Protection Training and the additional Data Protection for Research course prior to data collection (though I have completed this within 5 years). This study does not anticipate producing patentable or commercially exploitable outputs, and no collaboration arrangements with external commercial entities are planned; should this position change, Edinburgh Innovations will be contacted before any intellectual property agreements are entered into.

Storage and Back-Up

8) Where will your data be stored and backed-up during the project?

Stored on DataStore

Selection and Preservation

9) Where will the data be stored long-term?

At the conclusion of the project, the research dataset, comprising cleaned behavioural timing data, processed survey responses, coded qualitative data, and analysis scripts, will be deposited in the University of Edinburgh's Edinburgh DataShare repository for open-access, long-term preservation. DataShare is the preferred repository for this project as the dataset contains no sensitive or personally identifiable data; participants are identified by anonymous IDs only, and no further anonymisation steps will be required prior to deposit. All data will be accompanied by a README.txt file and sufficient metadata (including title, abstract, keywords, creator details, and funder information) to support discovery and reuse by other researchers. Deposition will take place at the point of publication or thesis submission, or earlier if required by the funding body. Raw qualitative transcripts or any unanticipated outputs that cannot be fully anonymised will be assessed on a case-by-case basis and, if necessary, deposited in the University's DataVault for closed, private preservation rather than open access. During the active project period, data will be shared with supervisors and collaborators via the University's approved collaboration tools, and the working dataset will be held on the University Research DataStore as described in the storage section above. No external disciplinary repository is planned at this stage, though this will be reviewed in light of any funder or journal requirements at the point of publication.

10) Which data will be retained long-term?

The following data will be retained long-term and deposited in Edinburgh DataShare at the close of the project: the cleaned and processed behavioural timing dataset, the cleaned survey dataset including the computed AI-use moderator composite, the coded qualitative dataset along with the codebook and inter-rater reliability outputs, and all analysis scripts.

Raw data files will also be retained and deposited alongside the processed versions in the interest of transparency, given that they contain no sensitive or personally identifiable information.

Data Sharing

11) Will the data produced from your project be made open?

- Yes: go to 12

12) How will you maximize data discoverability & access?

The dataset will be made openly available via Edinburgh DataShare upon deposit, which will assign a persistent Digital Object Identifier (DOI) to the dataset, ensuring it is consistently citable and discoverable over the long term.

If a relevant disciplinary repository (for example, one specialising in psychology or human-computer interaction datasets) is identified as more appropriate at the point of publication, deposition there will be considered as an alternative or complement to DataShare.

To further maximise discoverability, the DataShare deposit will be accompanied by comprehensive metadata: including title, abstract, keywords, creator and funder details. Links to associated publication will be made so that the dataset can be located through both the DataShare interface and wider academic search tools.

Responsibilities & Resources

14) Who will be responsible for the research data management of this project?

Jessica Ciupa
s1810129@ed.ac.uk

15) Will you require any training or resources to properly manage your research data throughout this project?

Yes.

Planned Research Outputs

Conference paper - "Automating the Hard Choices: Agentic AI's impact on Personhood and Moral Identity"

Conference paper for HAR: Human and Artificial Rationalities. A interdisciplinary conference on computer science, psychology and philosophy.

Planned research output details

Title	DOI	Type	Release date	Access level	Repository(ies)	File size	License	Metadata standard(s)	May contain sensitive data?	May contain PII?
Automating the Hard Choices: Agentic AI's impact o ...		Conference paper	2027-04-01	Restricted	None specified		Creative Commons Attribution 4.0 International	None specified	No	No