

Authors' response: Thank you to both reviewers for their comments. Our paper has been revised based on your suggestions, as addressed below.

Reviewer: 1

Comments to the Author

The paper has significantly improved. In taking this work forward, it will be important to keep in mind the specific barriers revealed by the interviews, as policy needs to be informed with sufficient granularity to be useful. This is a complex issue, and the legal questions must be engaged with.

Authors' response: Thank you for your comment and acknowledging the changes that we had made to improve the paper. To improve the engagement of legal questions as identified in your comment as well as the other reviewer's comments, we have clarified the positioning of the commons in context of existing legal frameworks and potential deployment. This is explained in more detail below.

Reviewer: 2

Comments to the Author

I think the problem is that currently the authors are trying to present their paper as both:

- 1. An empirical study analysing how a range of stakeholders apprehend / understand / assess data commons: this is valuable in itself, and does not require any legal argumentation*
- 2. A paper that should be taken as advice to policy-makers, which criticises the extent to which other data stewardship mechanisms are 'effective' when it comes to 'protecting' data subjects' personal data or 'helping them achieve their data protection aims'.*

The authors are right to point out that 1 does not require any rigorous engagement with legal arguments / legal structure of various data stewardship models. I do not think they can say the same about 2. The authors have improved the clarity and accuracy of what were previously very confusing statements, but key, important claims in the paper are left unsubstantiated. I think this is not ok given this paper's claim to be received as advice for policy makers.

This paper contains far too many inaccuracies for me to point them all out here. I have flagged some below -it takes a lot of time, and I do not think it is fair to expect reviewers to do this repeatedly- as it would take too long to pick them all out.)

When I first reviewed this paper I was taken by its potential, as the literature really does need papers on data commons, but this paper is still structurally reliant on far too many unsubstantiated (or plainly inaccurate) claims. To publish it in this state would prove actively detrimental to policy makers given the number of such inaccuracies or unsubstantiated claims. If, instead of merely mentioning (sometimes misleadingly) various sources the authors actively engaged with them and showed why X or Y is right or wrong on this or that point, it could be a much needed addition to the literature. Instead the authors tend to compile a list of sources assuming this backs their position when often it does not. In that

way it paints what I fear is a misleading picture of the data stewardship landscape. Given how crucial the latter is to their argument, I fear this paper must be rejected.

Authors' response: Thank you for your detailed comments and for taking the time to review our paper. We appreciate that you have acknowledged the difficulties of both conducting empirical work and presenting them within policy as this gap must be bridged in order to bring theoretical applications of data stewardship and data governance into practice. Following your recommendations, we have extensively reviewed the references included in our paper to ensure that they are in line with existing definitions and practices as supported by the cited authors' and their institutions. Where there is value in citing literature that may currently be considered outdated, we clearly note any differences.

As data stewardship and data governance are fast-moving areas, to ensure that our work reflects current research and attitudes, we have included new literature from organisations such as GPAI, The Data Economy Lab and Stanford among others, published between the revision submitted for the previous review (September 2021) and our current revisions (December 2021) that explore the benefits and challenges of adopting data trusts as well as other bottom-up data stewardship models. These have been included to clarify both the current state of the art as well as our viewpoints.

Unsubstantiated or inaccurate claims:

1. 'Data trusts could in theory support responses to certain data subject rights requests, particularly through access requests, but it may be difficult to benefit from other rights such as portability and erasure to support data subjects through trusts (Delacroix & Lawrence, 2019).'

Why? This needs to be explained.

Authors' response: This sentence has been expanded to clarify the potential tensions related to trade secrets and intellectual property which has been noted by the same referenced authors in the same paper and confirmed by the Ada Lovelace Institute.

2. 'While efforts have encouraged the creation of "bottom-up" data trusts that aim to empower data subjects to control their data in a way that acknowledges both the vulnerability and limited ability to engage with the day-to-day choices underlying data governance (Delacroix & Lawrence, 2019), many data trusts that have been created are top down in nature (Open Data Institute, 2019b).'

This is a misleading sentence. If the authors want to argue the above they need to explain how it is compatible with the following two facts:

- 1. None of the 2019 ODI 'data trusts' pilots relied on trust law.*
- 2. the Ada Lovelace report's clear, consensus definition (it was endorsed and co-produced by the Law Society and the ODI, thereby superseding the 2019 report -which is wrongly quoted*

earlier instead of the Ada Lovelace report- this shift was confirmed in several later ODI blogs), defines data trusts as relying on trust law.

If the authors prefer to refer to the ODI 2019 report because they disagree with the Ada Lovelace report's definition of data trusts they should say so and explain why, since the latter report was co-sponsored by the UK AI council, with direct links to the very policy makers this paper seeks to advice.

Authors' response: Thank you for pointing out the lack of clarity. The use of the ODI 2019 reference has been removed from any current definitions of data trusts. We believe that reference to the ODI 2019 pilot continues to be useful, as is also cited from many of the references we have included. In our paper, where the ODI 2019 pilot is demonstrated to improve the understanding on the benefits of data trusts, i.e. in supporting the initiation and use of data trusts, it is explicitly mentioned that trust law was not applied.

3. The authors criticise data coops because, among other things, 'cooperatives often rely on contract or incorporation to establish rights, obligations and governance, which could reintroduce some challenges related to collaboration and mobilisation the framework was intended to limit (Ada Lovelace 2021)'.

The authors cannot refer to this purported limitation of data coops without at the very least considering the extent to which the above challenge might apply just as much to the data commons framework they have in mind. This is actually a critical question, which will be of central importance to policy makers. The authors point out in their reply that they do not need to engage with legal points because it is not a legal paper, but their making the above point against data coops seems to introduce double standards here.

Authors' response: The question regarding how this challenge does and does not apply to data commons was mentioned in the 'A Data Commons for Data Protection'. However, based on your comment, it was evidently unclear. We have elaborated on how the commons differs, specifically in the legal context, within the same section.

4. On page 5, the authors state 'The use of data stewards and information fiduciaries (where similar legal safeguards are placed on the companies themselves as opposed to an independent trust) for the protection of personal data can further entrench existing companies in their data collection processes and fail to challenge those online platforms' structural power'

This is inaccurate. The authors refer to an American paper, which explicitly discusses the concept of information fiduciaries as it was put forward at some point in the US: the legal safeguards concomitant with information fiduciaries as delineated in that American proposal have hardly anything in common with the legal safeguards concomitant with trusts, say.

Authors' response: We agree that information fiduciaries are not the same as data trusts given that the fiduciary responsibility is different as we identified in the quoted sentence. However, the challenge we noted is one that can also apply to data trusts and other existing data stewardship frameworks, with impact beyond jurisdictions and borders. Given that the further exploration of whether data stewardship frameworks may entrench existing data

collection processes is beyond the scope of our paper, we have removed this sentence to avoid confusion.

5. There are no settlers in data coops

Authors' response: As we referred to different data stewardship mechanisms within the same sentence, our wording was unclear. This has been changed.

6. 'even if data stewardship frameworks are bottom-up, it is unclear how genuine and adequate engagement mechanisms can be deployed': this is a very important question, and it would be great if the paper addressed more concretely and precisely how the proposed data commons framework guarantees genuine and adequate engagement. What are the safeguards?

Authors' response: We agree that deploying genuine and adequate engagement is important, which is why we conducted interviews to understanding experts' views on data subject and wider stakeholder engagement. As outlined in our 'Discussion', based on our empirical work, we identify key considerations to be adopted as part of the data commons framework to ensure that appropriate mechanisms have been put in place to support engagement. These considerations have then been incorporated into the data commons checklist. To make this thread clearer, we have amended the 'Discussion' section to directly outline how engagement has been adopted within the commons as well as added a new 'Encouraging collaboration and data subject engagement in a commons' section to elaborate on how the commons applies existing public and user engagement methodologies in its framework.

7. The authors quote the ADa Lovelace report to claim that 'Data cooperatives often do not preserve privacy as a first priority'. Which part of the Ada Lovelace report exactly states this? Same question about 'advantages data subjects as the trust's beneficiaries may gain is not always clear': this needs to be unpacked and substantiated (many examples are given of potential benefits for the trust's beneficiaries, all dependent on examples: from insights about educational profiles / learning profile in the case of a school to insights about consumer behaviour / finance management / choice of healthcare providers etc).

Authors' response: To address concerns about the statements, the Ada Lovelace report states "Although privacy is usually a feature they respect, it is hard to find data cooperatives intending to preserve privacy as a first priority, through limiting the data that is collected and processed" and "Further work will be required to analyse the sorts of powers that a trustee tasked with stewarding those rights might be able to wield, and the advantages that might accrue to the trust's beneficiaries as a result". These quotes can be found in the references cited as well as the full 'Exploring legal mechanisms for data stewardship' report by the Ada Lovelace Institutes which collates the former references into one document.

With regards to unpacking and clarifying these statements, good examples of existing data stewardship models have been added to ensure balanced representation of real governance models. References to the Data Economy Lab's recent work that states the benefits and challenges to several data stewardship frameworks have also been added. We agree that it is important that the successes and real-world applications of different data stewardship models be accurately mentioned and taken into consideration when suggesting that a

commons could serve as an alternative means for data stewardship in context of data protection. We have included as much information as we see relevant to the wider body of research but would like to acknowledge that the clarification of all claims beyond reference to existing work would be a significant task and is one that has already been taken by multiple organisations cited in our paper. As our work focuses on gathering empirical data from existing commons and applying them more holistically to a policy framework that supports data subject engagement, we hope that the sources in which our research draws from accurately represents the data stewardship landscape while outlining how a commons can be useful in this area.

8. 'Existing data stewardship models may introduce trustees and experts that are able to prevent potential data-related harms (Ada Lovelace Institute, 2021b) but data subjects may still have to seek remedies from data protection harms such as breaches themselves, with limited support on how to exercise data subject rights under data protection regulations.'

That is, again, inaccurate. In the case of data trusts, a court oversees each trust, with broad intervention powers. It is also unclear why members of a data coop would necessarily be in a worse position than members of a data common in this respect: concerted / coordinated action is possible in both cases, and of course the traditional recourses are available in both cases. It would be important to unpack this, rather than state this as if obvious?

Authors' response: In an earlier version of our paper, it was pointed out from a reviewer that data trusts, cooperatives, and collaboratives do not support direct remedies related to remedying data breaches. We agreed with them and clarified our understanding in this area with reference to the statement quoted. However, there seems to be disagreement regarding the extent to which this is true as demonstrated by your comment. The clarification required here seems to be that supporting remedies is not mandatory within these frameworks. This is contrasted with the creation of a data protection-focused data commons that explicitly includes remedies through the exercise of data subject rights and the coordination of data rights efforts within the community, which we note throughout the article.

9. The comparative table is riddled with inaccuracies (some of them mentioned above). Apart from all the points above, the authors introduce all sorts of unsubstantiated distinctions. To only mention one example, why would it be the case that data trusts (or data coops) cannot provide 'a good governance model to minimise the risks of data breaches etc' (unlike data foundations)? This is simply stated, as if unproblematically true. In its present state, the table presents itself as a supposedly neutral, comparative analysis, when it is anything but (and arguably misleading).

Author's response: The table has been amended to include the most up-to-date references, including those published between the previous and current revision, and specific reference to data subject engagement.

It should be noted that all references and statements made in the table are summaries, albeit critical, of the reasoning that has been previously explored within the section of our paper. As noted throughout our paper and in many of the references cited, there is no one-size-fits-all solution to data stewardship and governance. The purpose of the table, the

inclusion of which was suggested by a previous reviewer, is to concisely summarise some of the arguments made within our paper as well as provide reference from which those claims are made. As a result, within the table, the benefit of one data stewardship model does not imply that it is not a benefit of another. This also applies to limitations. We do not claim that X model is 'the best governance model' for Y reasons, but only that research has been conducted to illustrate a model's benefits as well as limitations. In our table, we simply highlight, in what we believe to be accurate representations of the references included, how some data stewardship models may be more suited for different purposes.

I will stop here as I have simply run out of time, I could go on...I am sorry to have to be so negative on what I had hoped would be a much needed resource. I do believe rigour has never been more important to maintain and this paper simply does not pass a minimum standard given the arguments it purports to make.

Authors' response: Once again, thank you to all reviewers for their time and comments.