



Demonstrating and assessing trustworthiness when sharing data

Interim research report | September 2020



This is licensed under a Creative Commons
Attribution-ShareAlike 4.0 International license

Open Data Institute

Contents

Introduction	3
Initial findings	4
Trust and trustworthiness are highly context-dependent	4
Third-party certifications are useful, but only to a degree and in certain contexts	6
The most appropriate method will depend on the context	7
Plans for the development phase	10
Next steps	10

About

This report has been researched and produced by the Open Data Institute, and was published in September 2020. The lead authors are Jared Robert Keller, Sonia Duarte, Fionntán O'Donnell, Elea Himmelsbach, Renate Samson and Olivier Thereaux.

If you would like to send us feedback, please get in contact with us at research@theodi.org.



How can it be improved? We welcome suggestions from the community in the comments.

Introduction

When sharing or accessing data, organisations need to be able to trust those they are interacting with and feel comfortable doing so. Since data ecosystems are made up of a range of organisations with different roles and responsibilities, an organisation accessing or using data often needs to be able to trust more than the organisation holding or providing that data. To some degree they need to be able to trust any other organisations involved in the collection, storage, management, analysis, sharing or use of that data along the way.

At the Open Data Institute (ODI) we have been conducting a research and development project that aims to develop tools and resources to build trust between organisations when sharing and accessing data, with the goal of helping data ecosystems operate more effectively while reducing the risk of causing harm.

The scope of the project presents an interesting challenge, as it involves concepts that can be difficult to pin down. Our previous [research into designing trustworthy data institutions](#) has shown, for instance, that ‘trust’ and ‘trustworthiness’ are often perceived as nebulous topics.¹ Trust is inherently about relationships and communication between people and organisations, but it can be difficult for multiple parties to align around a shared understanding of what trust and trustworthiness mean. Similarly, sharing and increasing access to data can take many forms. Our research into the [wide world of data sharing](#) demonstrated that there is a range of approaches to sharing data, each of which involves different trust relationships and carries different benefits and limitations.² Lastly, [research by the ODI](#) and others found that although there is broad agreement that data sharing is beneficial, it is often difficult to quantify the actual ‘value’ or ‘impact’ of sharing data.³ The goal of this project – to *positively impact* data ecosystems by *building trust* between organisations when *sharing data* – is multidimensional and challenging, but worthwhile.

This interim report summarises the research conducted over the first phase of the project and looks ahead to the coming development phase. In particular, it lays out why third-party certifications and audits are useful, but only to an extent and in certain contexts. In other cases, alternative approaches to building trust may be more appropriate.

¹ Open Data Institute (2020), [‘Designing trustworthy data institutions’](#).

² Open Data Institute (2020), [‘Mapping the wide world of data sharing’](#).

³ Open Data Institute (2020), [‘The Value of Data’](#).

Initial findings

Given the necessarily broad terrain covered by this project, the research phase aimed to narrow the scope. We cannot address every challenge related to trust and data that organisations might have, so we have focused on defining what is in/out of bounds and setting achievable goals. We have outlined a few of our initial findings below.

An important learning to come out of our early research (backed up by [previous research at the ODI](#)) is that there is a difference between ‘being trustworthy’ and ‘being trusted’.⁴ The main difference is that being trusted is relational and involves an assessment by another party. An organisation might deem themselves *trustworthy*, but they can only consider themselves *trusted* once another organisation has placed its trust in them in some way – for instance trusting them to perform a service or deliver a product.

Stressing the relational aspect of trust is important, as it highlights that organisations need to demonstrate their trustworthiness to specific audiences rather than to the world in general. Indeed, our early research for this project has confirmed that some organisations think of trust and trustworthiness in general terms and do not think deeply enough about how to demonstrate specific aspects of their trustworthiness to specific audiences in specific ways.

In light of this, when we speak to organisations we ask them not about trust in general, but about how, specifically, they go about **improving** and **demonstrating** their trustworthiness to other organisations when accessing, using and sharing data – and vice versa, how they go about **assessing** the trustworthiness of others. For instance:

- An organisation might strive to **improve** its internal trustworthiness by undergoing a training scheme or through conducting ethical reviews.
- This is often not enough to garner the trust of other organisations, however. In order to be trusted, they must **demonstrate** their trustworthiness to other organisations, perhaps by displaying kitemarks or publishing the results of ethical reviews.
- And since ‘being trusted’ is relational, those other organisations must be able to **assess** the trustworthiness of that organisation, possibly by confirming that it has undergone a relevant training scheme or by locating openly-published results of ethical reviews.

⁴ Open Data Institute (2020), [‘Designing trustworthy data institutions’](#).

Trust and trustworthiness are highly context-dependent

Our research has confirmed that ‘being trustworthy’ means very different things to different people – and therefore the methods that people and organisations use to improve, demonstrate and assess trustworthiness vary widely depending on the context.

For instance, ‘trustworthiness’ means different things to different people and organisations.

- **Organisations working in different sectors.** What an organisation does to demonstrate trustworthiness in the health sector is different to what organisations do in the engineering or finance sectors. This is because different sectors have different rules governing ethics and trustworthiness, different governing bodies and different forms of sanction, redress and liability. In addition, different sectors value and use data in different ways, which can produce different business models and priorities within organisations.
- **Organisations performing different roles within a sector.** Even within the same sector, what defines a trustworthy ‘data contributor’ is often not the same as what defines a trustworthy ‘data intermediary’ or ‘data user’. Similarly, people and communities impacted by the sharing of data are likely to view trustworthiness differently than regulators, funders or policymakers working within that sector.
- **People working in different parts of an organisation.** Within a single organisation, a member of the legal department will have a different view of trustworthiness compared to someone working in data management, sales or communications, for example. One respondent to our survey, in fact, noted they were answering from only one perspective among many within their organisation: “This is in my role working with data, I am sure that if a different member of our organisation, for example, finance, completed this they would respond differently.”

Research and methodologies

Since May 2020, we have been conducting desk research on this topic, interviewing experts and surveying people from organisations involved in sharing or increasing access to data.

- **Expert interviews** – Between June and August we conducted 10 expert interviews with people from organisations performing different roles within data ecosystems – for example data contributors, data intermediaries and data users – across the health and finance sectors.
- **Qualitative and quantitative survey** – In August we launched a survey to identify how organisations demonstrate and assess trustworthiness when sharing data. The questions were aimed primarily at people working in organisations that collect, store, use, access or share data in some way, but we

welcomed responses from regulators, funders, policymakers and certifying bodies as well. To date, the survey has received 59 responses from organisations across a dozen sectors.

- **Working in the open** – Throughout this project we have strived to work in the open and have provided regular updates of our progress, including through the publication of an [introductory launch blog post](#), a [work note on the various research streams](#) of this project and a [work note detailing the research that informed the survey](#).⁵

The variables outlined above impact how trust and trustworthiness are defined, improved, demonstrated and assessed – but they are only a start. We have also identified other variables that impact trust and trustworthiness, such as:

- **The maturity of the use case or ecosystem**, for example an established ecosystem versus a new or evolving ecosystem.
- **The sensitivity of the data involved**, for example an agreement to share non-sensitive data versus an agreement to share highly sensitive or personal data.
- **The type of technologies in use**, especially newer technologies or technologies that develop or change quickly.
- **The type of service, product or transaction**, for example a business-to-business service versus a business-to-consumer service.
- **The role or importance of data within the organisation involved**, for example an organisation whose business model revolves around data in some way versus an organisation for whom data is a smaller part of their business.

Third-party certifications are useful, but only to a degree and in certain contexts

One consistent theme that emerged from our desk research, expert interviews and survey responses is that there is no single method, approach or mechanism that can build trust between organisations in every circumstance and every context. There is a wide range of different methods for improving, demonstrating and assessing trustworthiness, and different methods are better suited to different circumstances, challenges or contexts. What works to build trust between organisations in the health sector might not work in finance; and what works to build trust between two organisations when sharing data directly might not work for an ecosystem of organisations sharing data via an intermediary.

Our research has found that, in particular, while third-party certifications and audits are useful for improving, demonstrating and assessing trustworthiness, in some cases alternative approaches may be more appropriate.

⁵ The Open Data Institute (2020), '[Help us understand how certification can help build trust in data ecosystems](#)'; The Open Data Institute (2020), '[R&D: Building trust through audit and certification. Worknote #1](#)'. The Open Data Institute (2020), '[R&D: Building trust through audit and certification – Worknote #2](#)'.

For instance, when we asked respondents to rate the usefulness of things like third-party assessments, audits and certifications in helping them demonstrate their trustworthiness to others, the average rating was 7.8 out of 10 (10 being 'very useful'). Similarly, 71% of respondents said they were either satisfied or very satisfied with the certification schemes that existed in their sectors. This was borne out in many of our interviews, with one interviewee noting that third-party assessments are often needed to "reduce the risk" of "bad things happening".

However, certifications and audits are often seen as only one approach to improving, demonstrating and assessing trustworthiness. Indeed, one respondent to our survey wrote that while third-party assessments were "of some use", they were ultimately "only an indicator not a guarantee". This came out in our interviews as well, with some experts noting that in order to build trust between organisations, a range of different approaches often need to be deployed in tandem. Certifications and audits were seen as offering a useful starting point or foundation, but because trust and trustworthiness are highly context-dependent, other methods are often necessary, depending on the situation. As one interviewee put it: "Certification could provide a very baseline level of assurance about a data source, but it would just be the baseline because the data that you need is very situational, specific."

In some circumstances or contexts, certifications and audits are seen as unproductive or unsuitable, for a range of reasons, including:

- A belief that certification or audit processes often cannot keep up with the pace of change in the domain, especially the pace of technological change.
- A concern that some assessment schemes can become 'tick-box exercises' that lack robustness.
- A concern that some certification and auditing schemes are easily gamed or become 'certification theatre', where both parties know the certification or audit is inadequate but have no incentive to improve it.
- A feeling that some complex technologies or data flows, for instance machine learning algorithms or a system with millions of new data points daily, may be hard, if not impossible, to adequately assess.

Finally, some of our interviewees and respondents noted that third-party assessments can, in certain cases, actually be harmful. In particular they noted that certification and auditing schemes can hinder innovation if they are costly and time consuming. This was seen to lock some smaller organisations out of the market, especially if the schemes are mandatory. One of the respondents to our survey, for instance, noted: "As a startup, I am concerned about losing access due to higher 'costs of doing business', or worse, a 'pay-to-play' culture."

The most appropriate method will depend on the context

Because third-party assessments, certifications and audits are only useful in some contexts, and because there is a wide range of different methods for improving, demonstrating and assessing trustworthiness, we surveyed a range of organisations in different sectors to gather their views on the value of third-party assessments and to understand how they demonstrate and assess trustworthiness when sharing data.

The first section of the survey focused on identifying which aspects of an organisation are the most important when it comes to demonstrating trustworthiness. Organisations are multifaceted, after all, so demonstrating trustworthiness is multifaceted as well. By drawing on our desk research, expert interviews and other research at the ODI we have been able to put together a list of different aspects of an organisation that organisations might feel the need to show are trustworthy. In the survey we posed the question, ‘what about your organisation do you try to demonstrate is trustworthy?’ and asked respondents to rank the list in terms of priority for their organisation.

The list currently includes things like:

1. Your legal structure and compliance with relevant laws or regulations
2. Your software and technical infrastructure
3. Your information governance framework (eg data management and data protection procedures)
4. Your commercial practices and how ‘value’ is created or shared
5. Your decision-making processes for how data is collected, managed, used and shared
6. Your ethical review procedures and how you work to minimise bias or harms
7. Your oversight, accountability and redress procedures
8. Your transparency or engagement procedures
9. The capabilities and expertise of the individuals working within your organisation (for example demonstrating that they are good practitioners)
10. The datasets you supply (for example quality, safety, value)
11. The services you deliver (for example quality, safety, value)
12. The products you provide (for example quality, safety, value)
13. Your commitment to ‘public good’ or having a ‘positive impact’
14. Your diversity and inclusion practices
15. Your financial sustainability
16. Your environmental sustainability

We asked the same question for *assessing* the trustworthiness of other organisations. The list is a work in progress, so if there is anything that you believe is missing from our list, please [get in touch](#) to let us know.

A few noteworthy initial findings

- **Respondents to our survey stressed the importance of being able to demonstrate their commitment to ‘public good’.** When asked ‘what about your organisation do you try to prove is trustworthy?’, a commitment to ‘public good’ or ‘social impact’ was ranked the second most important aspect of an organisation. However, when asked ‘which aspects of other organisations are important to assess?’, a commitment to ‘public good’ or ‘social impact’ was ranked eighth. More work will need to be done to identify how representative the results of our survey are, but they suggest that while organisations are often concerned with *demonstrating* their commitment to public good, many appear to be less concerned with *assessing* whether others share that same commitment.
- **Efforts to build trust have frequently been reactive rather than proactive.** Throughout our research we have come across instances where methods have been developed for improving, demonstrating or assessing trustworthiness as a reaction to an event, such as a reputational issue, a legal challenge or a system security concern. We also found that unless there is a clear incentive, developing methods for building trust can be drawn out, complex or difficult to agree. However, a starting point for organisations can be first engaging in low-effort or low-risk situations, where trust is much more easily built. In these cases, it is less about demonstrating trustworthiness and more about having no reason to distrust. From these situations, trust can be built for more complex or higher-risk situations.

In addition to asking people *what* about their organisation they strive to demonstrate is trustworthy, the survey also asked people *how* they go about doing so. We wanted to identify to what extent organisations use other methods or approaches to demonstrate and assess trustworthiness outside of third-party assessments, certifications and audits. Through our desk research, interviews and survey we have assembled a list of almost 20 ways that organisations demonstrate trustworthiness, including:

1. Committing to specific legal agreements – for example contracts, terms and conditions or licenses
2. Abiding by relevant laws and regulations
3. Abiding by relevant standards (technical or data or quality)
4. Committing to relevant principles or values
5. Committing to relevant best practice or codes of conduct
6. Committing to relevant organisational or industry norms
7. Attaining relevant certifications (at an organisational level)
8. Attaining relevant certifications (at an individual level – that is, for practitioners within the organisation)
9. Displaying relevant kitemarks, stamps or labels
10. Undergoing tests or exams (either of your organisation, your people, or your products)
11. Undergoing or undertaking auditing schemes
12. Answering queries or sitting for interviews
13. Agreeing to performance monitoring or oversight

14. Agreeing to penalties or redress mechanisms
15. Embedding regular internal review processes
16. Embedding independent advisory processes – for example citizen juries, ethics panels or data access boards
17. Committing to proactive transparency – for example the publication of earnings reports, decision logs or the minutes of meetings
18. Engaging and communicating with users, stakeholders or community members
19. Through implementing technologies designed to underpin or improve trust

As above, if there is anything that you believe is missing from our list, please [contact us](#) to let us know.

Plans for the development phase

Providing an organisation with an approach to how they can build trust depends on a wide range of variables. Therefore the development phase of this project will focus on working with people and organisations in a range of contexts to identify their specific trust-related challenges and potentially co-create guidance, tools or resources that can help them address their specific needs.

Focusing on addressing concrete, trust-related needs will enable us to take an in-depth look at real data ecosystems, the trust relationships between different actors and how trustworthiness is improved, demonstrated and assessed within those ecosystems. Furthermore, by creating, testing and iterating potential tools or resources with the actors and stakeholders involved in those contexts, we can help ensure that what we produce will be fit-for-purpose, useful and therefore more likely to be adopted. A useful approach is unlikely to be developed unless we are led by the people who understand the domain; our research has shown that creating a certification or audit that is inadequate can actually be harmful since, among other things, it can give organisations a false sense of security and leave them unaware of potential risks. Our research has also shown that there must be transparency around what is being certified, who is doing the certifying and why. As one of the respondents to our survey put it: “If I have no real say in what the standards are, I'm much less likely to buy into them.” One of the best ways of helping to ensure this is to develop them with input from the relevant communities.

In some cases, we may find that people and organisations in a specific context would benefit from an auditing scheme or agreed standards that could eventually form the basis of a certification scheme. But in other cases we may find that things like guidance, training tools or a checklist of resources would help organisations improve, demonstrate and assess trustworthiness more effectively than standards and certifications.

The insights and lessons we draw from this development work will be applicable beyond the specific ecosystems and use cases we looked into. We therefore see the development phase as ultimately contributing to our ability to offer overarching guidance and lessons to organisations in disparate contexts.

Because of the nature of this work, we see it as contributing to a necessarily longer term research endeavour. Further work from ourselves or other interested organisations will undoubtedly be needed.

Next steps

Over the next six months, this project will progress through discovery, to alpha and then beta. During the discovery phase we will identify pressing challenges or opportunities related to trust and data, test some of our initial findings with organisations in those areas, identify their specific trust-related challenges and develop initial ideas for tools or resources that could address those challenges. As we move into the alpha phase, we will develop basic prototypes that we can put in front of potential users. The beta phase will involve taking the best idea(s) from the alpha phase and creating a refined version, ready to test on users.

Based on our research, one of the areas where we believe we can have the greatest impact in improving trust between organisations is in ‘less mature’ contexts. Our research has shown that in mature, traditional or static contexts, organisations tend to have well-established mechanisms for improving, demonstrating and assessing trustworthiness; in less mature contexts, previous understandings of trust and trustworthiness may need to be renegotiated and redefined and previously satisfactory methods for building trust may need to be recontextualised, updated or replaced. In these types of areas, building trust is particularly important, but particularly difficult.

Within the development phase, we are therefore interested in exploring a range of different ‘emerging’ contexts, such as:

- **‘New’ use cases** – for example, new uses of data or the repurposing of data in novel ways. For instance, the mobility and transport sector where location data is increasingly being repurposed for novel uses; or the health sector where data not traditionally viewed as health data (eg fitness data) is being used to make health decisions. These novel uses often outpace the development and adaptation of methods previously used to build trust when sharing data, so it can be difficult for organisations to improve, demonstrate or assess trustworthiness around these new uses of data.
- **‘Evolving’ data ecosystems**, for example, ecosystems where new intermediaries are being installed or organisations are taking on new roles as data stewards or data institutions. For instance, the health or mobility sectors where organisations are seeking to play the role of independent, third-party stewards of data. The introduction of these ‘new actors’ into existing ecosystems requires the existing actors in those ecosystems to reassess their roles and to potentially adapt existing mechanisms (or introduce new mechanisms) for improving, demonstrating or assessing trustworthiness.
- **‘Emerging’ sectors**, for example, less-established sectors where roles and responsibilities have yet to be settled and methods of building trust have yet to be agreed. For instance, the new sector that is consolidating around ‘alternative’ forms of investment data, or ‘alt data’. Within this new sector, data is being used in new ways, new data contributors are proliferating and organisations are unsure who they can trust. These emerging sectors often have ad hoc ways of improving, demonstrating or

assessing trustworthiness, but few established, scalable ways of doing so and little shared understanding of best practice or standards.

We will also be working with Frontier Economics, an economics research firm, to quantify the economic impact of increased trust between organisations in data ecosystems. At the ODI we have conducted research into the ways that trust can be established and built when sharing and accessing data. There is also a large amount of research into the relationship between data sharing and economic impact. The research with Frontier Economics will seek to address the evidence gap between these two areas and attempt to quantify the economic impact of increased trust between organisations when sharing and accessing data.

If you are an organisation that collects, manages, shares, accesses or uses data, [we would love to speak with you](#) about how you improve and demonstrate your trustworthiness and assess the trustworthiness of others.

During the development phase of this project we are likely to convene stakeholder workshops to test initial designs and gather feedback. If you would like to take part in these workshops or provide feedback as a 'critical friend' during this phase, please [get in touch](#).

We believe, after all, that only by collaborating with stakeholders and members of data ecosystems will we be able to help organisations build trust when sharing and accessing data, thereby helping to increase the sustainability and effectiveness of data ecosystems while reducing the risk of causing harm.