Data Trusts

Kieron O'Hara*

Recent years have seen the burgeoning of a literature on *data trusts*, and the unwary might therefore be led to believe that it is an idea whose time has come. Unfortunately, the ideas of the various authors who have contributed to this literature, who include the present author¹, haven't always coincided, and have been aimed at different problems at different levels of detail and hand-waving. We might therefore say more accurately that 'data trust' is a *brand* whose time has come, which in itself is a not uninteresting phenomenon, worthy of consideration.

The general background to this flourishing body of work is, of course, the big data revolution, the benefits it promises and the problems it is causing, as technologists, lawyers, individuals (either as citizens or consumers), governments, businesspeople, scientists, social scientists, medical scientists, engineers, designers of games, propagandists, advertisers, secret policemen and Russian troll farmers adjust to and get to grips with it. The combination of opportunity and threat is important – data trusts are pitched by their champions as means both of combating a threat and enabling more effective exploitations of opportunities.

The term 'data trust' trades upon the image of a medieval legal instrument, the trust, in which property is managed on *fiduciary* lines, so that the *trustee* runs it in the interests of the *beneficiary*. The idea was present in Roman Law, but reached fruition in common law jurisdictions, most obviously England. Yet it is not an instrument of common law. Its roots are in equity, the body of judgment of the Court of Chancery specifically designed to address injustice arising from strict application of the law. Trusts nowadays may be encountered in the administration of property inherited by minors, or in the management of commonly-held assets. The point of the fiduciary arrangement is the recognition that different and even opposing interests are in play and need to be balanced.²

The idea of a data trust speaks to a general notion of unease in the rich democracies that the data-driven world is out of our control, not being run for our benefit, and exploiting us in new and increasingly devious ways, the sort of unease articulated by such commentators as Zuboff³, who excoriates surveillance capitalism, and Ekbja and

DOI: 10.21552/edpl/2020/4/4

* Kieron O'Hara, Web and Internet Science Group, Electronics and Computer Science, University of Southampton, United Kingdom. For Correspondence: <kmoh@soton.ac.uk>.

1 Kieron O'Hara, 'Data Trusts: Ethics, Architecture and Governance for Trustworthy Data Stewardship' (2019) Web Science Institute White Paper, 1 <<http://dx.doi.org/10.5258/SOTON/WSI-WP001>>.

2 J.E. Penner, *The Law of Trusts* (Oxford University Press 2016).

3 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (Profile 2019).

Nardi⁴, who rail against the appropriation and use of data voluntarily created by us. As individuals, our data is of little value, but when data about billions of episodes in the lives of hundreds of millions of individuals is aggregated, the whole produces eye-wateringly valuable insights. New technology is expanding the range of data gathering all the time, and the global pandemic has even helped, by pushing more interactions online, creating new and richer data trails. The unease is not going to go away.

This is perceived as an issue for three sets of people, all of whom have reached for the notion of a data trust as the solution, but for very different problems: the first wants to increase the amount of data *gathered*, the second wants to increase the amount of data *shared*, while the third wants to *decrease* the amount of data shared.

Data Trusts for Data Aggregators

First, there are the aggregators of this data. Rich as Croesus, they intend to remain that way. In the West, these include companies of which we are all aware – Alphabet, Amazon or Facebook – together with more niche players such as Netflix or Palantir. The risk such companies face is that large aggregations of data may become the target of antitrust investigations, and be broken up, if they cease to be socially acceptable.

For these players, a data trust is a comforting-sounding name which conveys to a sceptical audience an arms-length engagement with the data created by a project. In a prominent example, Alphabet's subsidiary Sidewalk Labs, as part of its proposal to build a smart city on a waterfront development in Toronto, suggested an Urban Data Trust (UDT) to manage the sensitive data which would be generated.⁵ The UDT was a means of data governance which would acknowledge the range of interests to be accommodated, and reassure opponents that Sidewalk Toronto was not merely an attempt to add more data to Alphabet's hoard.

The UDT raised a number of issues. Who could use it to access the data? How open was it to be, and what access controls would be in place? Would access be in real time? There was a centralisation problem: if Sidewalk had a monopoly on data-collecting smart city apps, then developers would be aggrieved, but if the developers were free to join in, then Sidewalk couldn't guarantee that data collection would be fair and privacy-preserving. This meant that the UDT had to include a code of practice. Meanwhile, concerns about privacy and data sovereignty, and the USA PATRIOT Act in particular, meant that the UDT would have not only to manage the data, but also to store it in Canada. Its design had to address all of these concerns.

The UDT was not a literal trust, as defined in trust law. Sidewalk was keen to promote its evolving data governance plans as innovative and unprecedented, but flesh never

4 Hamid R Ekbia and Bonnie A Nardi, 'Heteromation: And Other Stories of Computing and Capitalism' (M.I.T. Press 2017)

5 Teresa Scassa, 'Designing data governance for data sharing: lessons from Sidewalk Toronto' (2020) *Technology and Regulation*, 44–56.

appeared on the bones. The UDT was neither a trust, nor a public sector body, nor a commons, nor a cooperative, nor a new structure requiring legislation. It was to be designed by Sidewalk, not by the subjects or the users of the data. Whatever it would have been, it was surely unlikely to create trust where it was lacking.

Thanks to the COVID pandemic, the project died, and the UDT with it. The UDT itself ended up less a coherent design than a series of *ad hoc* responses to public concerns. One specific difficulty was that the project as a whole was mistrusted, not merely the data-handling aspects of it, and the UDT was sometimes invoked to reassure those worried about the instrumentation of their environment *per se*, independently of what happened to the data afterwards. The top-down development of the concept by Sidewalk Labs, no doubt sincerely intended to anticipate and address public concern, simply underlined the broad feeling among opponents that this was a project out of their control. Even if they had access to the data produced, they could not influence the governance⁶.

Functional Data Trusts

The second set of people interested in data trusts are those who would like to aggregate data but, for whatever reason, can't. This usage of the term was pioneered by Hall and Pesenti's report on the UK AI industry⁷, which noted that the promise of the AI industry requires plenty of data analysis, and plenty of data sharing by those companies which, unlike Alphabet and others, are not blessed with giant troves. Yet, as they observed, it is harder to share data – even non-personal data – than one might think. Even those with a common interest in sharing data can struggle to get sign-off, or to write the contract that enables the share.

There are myriad reasons why this might be so. The receiver might worry about quality. The sharer might be concerned that others will make profits out of its IP, or that it will suffer reputational damage from misuse. Potential partners might be simply unaware of the existence of the data, or the companies, that would help them. The data may be commercially or personally very sensitive, making managers risk averse. There may be cross-jurisdictional complexities. This idea of a data trust, then, would be to create some kind of institution to bridge these gaps between sharers and receivers.

Hall and Pesenti are vague about what these data trusts would be, and this is inevitable, given that whatever prevents data sharing in any particular case will be unique to that context. For that reason, in my own work – which ploughs this particular furrow – I suggested that data trusts could be defined not as institutions, but *functionally*, as whatever would create the trust between receiver and sharer to enable data to be shared

⁶ *ibid*

⁷ Wendy Hall and Jérôme Pesenti, 'Growing the Artificial Intelligence Industry in the UK' (2017) Department for Digital, Culture, Media & Sport and Department for Business, Energy & Industrial Strategy <<https://www.gov.uk/government/publications/growing-the-artificial-intelligence-industry-in-the-uk>>

where it previously was not. This might be a legal device, most obviously a contract, but might be another kind of agreement – a code of conduct, or a computational architecture that might use cybersecurity measures to prevent undesirable uses taking place, or blockchain for auditing.⁸

The problem of how to increase the amount of data shared, whether for profit or science, was also taken up in a series of projects by the UK's Open Data Institute (ODI), which experimented with a number of structures that supported the sharing of sensitive data, and these provide some of the best examples of what they might be⁹. Joint projects between the ODI and the University of Southampton explored the symbiotic cooperation of large data-holders with small data-hungry startups¹⁰. Hall herself pioneered the idea of a Web Observatory, where Web-based data, together with visualisation and analytic tools, may be accessed¹¹.

These institutions are all envisaged to curate (not necessarily store) data while recognising the range of interests respecting its use. As I noted in my own contribution, these institutions are highly unlikely to be literal trusts in respect of trust law¹². The reason is quite straightforward. The object of a trust – the data in this case – is managed by the trustees for the sole benefit of the beneficiaries. Consider someone we might call DC – by this I mean a real data controller in the case of personal data, or someone in the analogous position of deciding the means and purpose of processing non-personal data. DC may want to share the data, but is he or she likely to be keen on surrendering control to the trustee, and thereby giving that trustee *carte blanche*, even if DC is named as one of the beneficiaries? One potential solution is for DC to be the trustee as well as a beneficiary, but this would fail in a literal trust as the trustee shouldn't be working for his or her own benefit.

The position of 'trustee', therefore, can only be metaphorical in a functional data trust, as would be the fiduciary duties. This kind of data trust, if well-designed, would ensure, either by technological means, or legal means, or an ethical code of conduct, or social pressures, that the data was used for the benefit of the beneficiaries (and possibly of others too, e.g. for social benefit). Outside of trust law, DC could be at once a trustee, a beneficiary and the determiner of the means and purpose of processing, ensuring that he or she shared the benefits without relinquishing control. Such an arrangement would be designed to ease the doubts of all parties to the share, and build confidence for a longer-term relationship, without the risk posed by personal liability for breaches. Sanctions would be engineered within the trust – for instance, ejection from

8 O'Hara (n 1).

9 Jack Hardinges, 'Data Trusts in 2020' (2020) Open Data Institute <<https://theodi.org/article/data-trusts-in-2020/>>

10 Sophie Stalla-Bourdillon, Gefion Thuermer, Johanna Walker, Laura Carmichael and Elena Simperl, 'Data protection by design: building the foundations of trustworthy data sharing' (2020) 2 Data and Policy 4.

11 Thanassis Tiropanis, Wendy Hall, Nigel Shadbolt, David De Roure, Noshir Contractor and Jim Hendler, 'The Web Science Observatory' (2013) 28 IEEE Intelligent Systems 2, 100-104.

12 O'Hara (n 1).

the trust – as part of its terms and conditions of use, rather than the liabilities of trust law.

Bottom up Data Trusts

The third group who might benefit from a data trust are data subjects themselves, who might contrive to place their personal data in a literal trust, to be managed for their benefit, or so that they at least get a cut of any profits made from them. This idea was first mooted in a consumer protection setting¹³, and then was adapted to address the power asymmetries between data controllers and data subjects that many believe are not adequately redressed by individual data protection rights (the *Facebook problem*), by Lawrence¹³. This approach has been defended most rigorously by Delacroix and Lawrence¹⁴.

Delacroix and Lawrence envisage an ecosystem of ‘bottom up’ data trusts, where trusts compete with each other to provide services, and data subjects choose where to place their data according to the benefits they wish to prioritise. Data consumers would negotiate with the trusts for access, and the trusts must negotiate acceptable terms for their members. As these are literal trusts in law, the trustees’ undivided loyalty must be to the beneficiaries (the data subjects whose data they manage), and so they would be liable for a breach if they ‘sold out’ on terms worse than anticipated. Trusts have traditionally been associated with property, but Delacroix and Lawrence argue convincingly that they can apply to rights over data use.

The power asymmetry, on this analysis, is down to data subjects acting as individuals (for example, using their data protection rights). Bottom-up data trusts correct the asymmetry by pooling the data resources of participating data subjects, to provide economies of scale, as well as creating a large data resource to be managed by the trustees, a valuable asset for which the trustees could demand concessions and benefits for their beneficiaries from data consumers.

The approach is attractive, and has the merit of a precise proposal respecting existing law to address a well-understood problem. However, there are questions about its practicality, and the business models upon which it would run.

The Settlor Problem

In the first place, we must ask who the settlor would be. It is unlikely to be the data controller. To see this, suppose a controller wants to share some personal data for profit and public good, in such a way as to have no direct impact on the data subjects, for example to train a medical ML algorithm. If GDPR (or the law appropriate for their ju-

13 Neil Lawrence, ‘Data trusts could allay our privacy fears’ *The Guardian* (London, 3 June 2016)

14 Sylvie Delacroix and Neil D Lawrence, ‘Bottom-up data trusts: disturbing the “one size fits all” approach to data governance’ (2019) 9 *International Data Privacy Law* 4, 236-252.

risdiction) allows the share to take place, then the controller has no incentive to form a data trust – he or she can simply retain control and share the data. If GDPR forbids the share, then it is hard to see how forming a data trust will enable it to take place. The controller would hand over the data to a trustee, whose loyalty has to be to the data subjects. The controller's interests drop out of the equation (as does the public good), and since the data subjects gain nothing from the proposed project, the trustee has no mandate to proceed. Hence either way, there is no incentive for the data controller to set up a data trust.

The settlors, then, must be the data subjects, using their GDPR portability rights to gain access to the data and delegate their rights to a trustee. This accords with Delacroix and Lawrence's use of the term 'bottom up.' But we have a dilemma. The *raison d'être* of this system is the power and information asymmetry between subjects and controllers. Subjects struggle to provide informed consent to the use of their data and thus of managing their rights¹⁵. Hence they put it in the hands of trustees to manage those rights for their benefit. However, trust law also demands 'certainty of intention' to create legal relations – the settlor must really have wanted to create a trust. Equity looks to intent, not form¹⁶. It is not just signing over rights – the settlor must understand precisely what is entailed by the trust context.

Now the dilemma is evident. The model assumes that the data subject is *unable* to give informed consent to a data controller, or to understand the implied contract in the privacy policy, but at the same time is *able* to understand and initiate legal relations with the trustee. The former may be described with a policy as long as *Hamlet*, but then so may the latter be. There is an analogous information asymmetry between data subject and trustee, and between data subject and controller, and it is not immediately obvious why the second asymmetry is ground for doubting that consent is informed, while the first is unproblematic.

The Trustee Problem

The second problem is that there are few obvious incentives to attract talented trustees. Failure to pursue the interests of the beneficiaries will render them liable for breach of trust, so the job itself is a risk. The trust's business model must provide for its running costs, the expenses of the trustees, and possibly a salary. The beneficiaries could be charged for the service, but they would surely expect sufficient benefits so that they are not making a loss by joining the trust; a subscription is therefore unlikely to be the only income stream. That means that the trust will have to charge data consumers for access to the data. This creates a dilemma. If the trust was set up to preserve privacy and only share the data sparingly, it would generate a correspondingly low income,

15 Solon Barocas and Helen Nissenbaum, 'On notice: the trouble with notice and consent' (2009) in Proceedings of the Engaging Data Forum: The First International Forum on the Application and Management of Personal Electronic Information <https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2567409>

16 J.E. Penner (n 2) 109 - 203.

unlikely to fund efficient management. To generate a large income, access to the data would have to be licensed to many consumers, and so the data subjects' privacy would be relatively compromised.

The danger is that we end up with a possibly extremely valuable trove of data managed unprofessionally and incompetently by underpaid and under-incentivised trustees. One imagines Mark Zuckerberg licking his lips at the prospect.

The Ecosystem Problem

Thirdly, Delacroix and Lawrence envisage an 'ecosystem' of bottom-up data trusts competing for data. This requires that we solve the above problem, so that there are enough trustees to populate an ecosystem. But the problem is wider; to paraphrase Shakespeare, you can call ecosystems from the vasty deep, but will they come when you call? The history of data protection is littered with user-centric schemes to empower data subjects, from P3P to personal data stores/clouds/vaults to personal information management to pods. The current author indeed has contributed to the literature¹⁷ (Van Kleek & O'Hara 2014), which tends to emphasise widespread take-up, tipping points, critical mass, and ecosystems. After two decades of research, there is no evidence that, outside the academic-legal complex, any demand exists among data subjects either to manage their own data, or to hand it over to a proxy.

The Facebook Problem *redux*

Fourthly, suppose all the problems are solved and we have our ecosystem. Data trusts compete, both as suppliers and as custodians of data. The wider their beneficiary base, the larger the datasets they can make available to consumers, so however they balance their income streams, success will be reflected by size. But this means that the data trust ecosystem is as subject to network effects as the current data market. As a trust gets larger, it will be in a better negotiating position vis-à-vis consumers, and will be able to demand greater benefits for its beneficiaries. Data subjects will be attracted to larger trusts, and smaller ones will tend to fade away, leading to oligopoly.

Of course, data subjects may have other benefits than financial ones in mind (they may value privacy and not mind that their data is not used much). However, the trustee problem remains for the privacy-respecting data trusts; if they are not selling much data, then they must generate income to cover their costs. Only the beneficiaries can contribute, and few are prepared to pay for privacy.

The result of solving the settlor problem, the trustee problem and the ecosystem problem, then, is likely to be a small number of behemoths with vast market power based

¹⁷ Max Van Kleek and Kieron O'Hara, 'The future of social is personal: the potential of the Personal Data Store' in Daniele Miorandi, Vincenzo Maltese, Michael Rovatsos, Anton Nijholt and James Stewart (eds.), *Social Collective Intelligence: Combining the Powers of Humans and Machines to Build a Smarter Society* (Springer 2014) 125-158.

on network effects – but that is the problem that bottom up data trusts were posited to solve in the first place.

Conclusion

We might conclude, therefore, that literal data trusts aren't going to solve the Facebook problem, and neither will metaphorical data trusts work as PR exercises for the tech giants. The data trusts with the greatest prospect for success are functional data trusts to encourage data sharing. However, against the charge that the notion of a functional data trust is so vague as to be virtually meaningless, I'm not sure I have a particularly compelling reply. The reasons that mutually beneficial sharing may break down are many and varied, and if bespoke solutions will always be needed for essentially *sui generis* problems, then functional data trusts may not be practicable either as a general solution. Nevertheless, such solutions are the best placed to emerge.