moz://a

Failure Modes for Data Stewardship

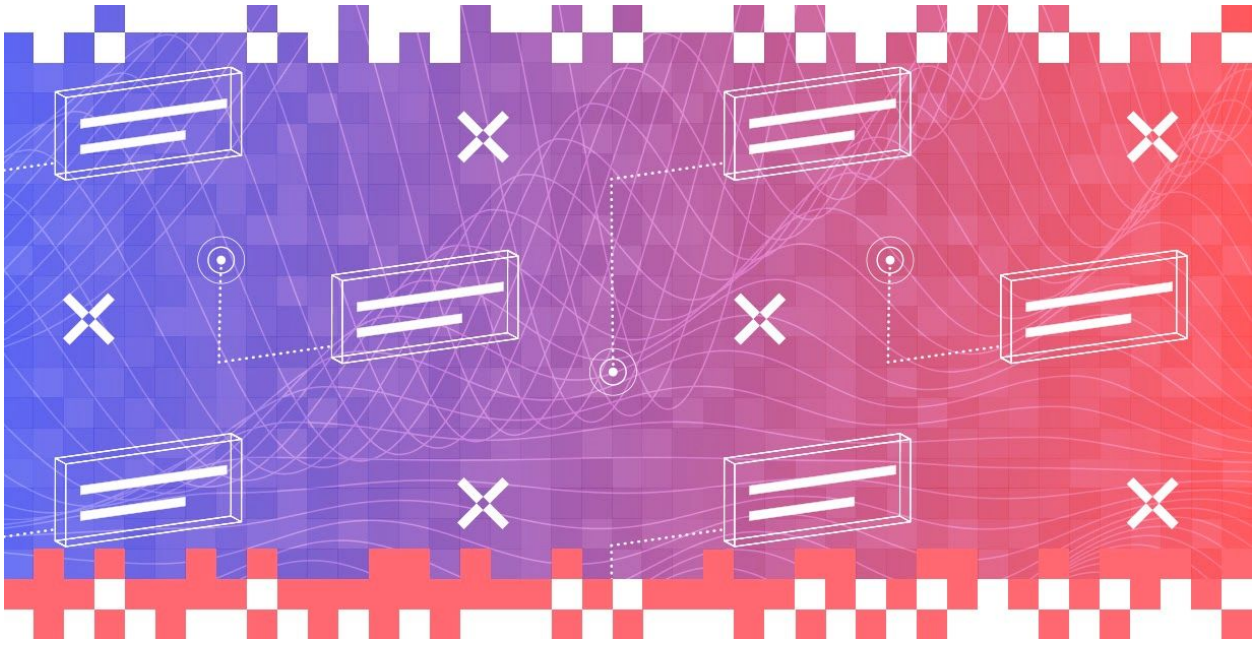

Author: Keith Porcaro for Mozilla Insights

September 2020

A note from Mozilla Insights

In the research series on data stewardship, Mozilla Insights assumes that there could be the opportunity to shift power imbalances through responsible innovation on data governance. However, we also recognize that any data initiative is accompanied by serious risks and hurdles. To keep our feet firmly planted on the ground, we commissioned Keith Porcaro to imagine theoretical scenarios, inspired by reality, where data initiatives may fail or go wrong.

Table of contents

Introduction

This brief examines failure modes for data stewardship initiatives. It was commissioned as a "red team" brief by Mozilla Insights, to help Mozilla and its partners improve their strategic planning, and determine how to support the data stewardship field.

The heart of this brief is a series of speculative narratives: stories about hypothetical future projects that have failed, and why. Although this brief looks at the data stewardship field with a critical eye, these stories are not destinies. While no single document can comprehensively capture all of the ways a data stewardship initiative might fail, this brief succeeds if it helps spark the imagination, and prompt thoughtful approaches to mitigating the harm to people that failure can cause.

True to the red team form, this brief was drafted in isolation from the rest of the research outlined here, and does not represent the opinions or positions of Mozilla or any of its staff, affiliates, or partners. The stories are fictional, and should not be read as descriptions of any current or past initiatives. Mistakes and omissions are my own.

How this brief is organized

In commissioning this brief, Mozilla Insights defined success for their data stewardship strategy as:

> *We must **reimagine** new models of managing data in the public interest; **reconstitute** collaborative spaces of design, development and deployment of new methods of data stewardship so they are more inclusive; and ultimately we must **rebalance** the power dynamics of our digitized worlds with counterweights.*

This brief's narratives are organized into three failure parts; inversions of Mozilla's provided success measures:

1. **Inertia.** New norms, initiatives, and technologies fail to take root, or fail to give individuals and communities sufficient protection or agency.
2. **Exclusion.** Data and technology projects continue to exclude, discriminate, and marginalize. New data governance models create new barriers to meaningful participation, or reinforce old barriers.
3. **Entrenchment.** Power asymmetries persist and grow worse. Entrenched interests co-opt, outcompete, or crush alternative models.

Each part contains several speculative failure narratives, matched with general descriptions of the failure implicated. Each part concludes with links to further reading that help illuminate the challenges that data stewardship initiatives face.

Part I: Inertia

New norms, initiatives, and technologies fail to take root, or fail to give individuals and communities sufficient protection or agency.

This part explores why individual data stewardship initiatives may fail, in the traditional sense: they don't attract users; they lack the power to achieve their goals; they are legally or financially untenable; they are conflict-riven; and so on. Many of these failure cases are not special: most organizations are susceptible to them, too. But that should not diminish their importance. A data stewardship initiative invites people to depend on them, and to trust them. The learnings that are extracted from an initiative's failure are ultimately of little use to the people who are let down or harmed by it.

Data stewardship initiatives fail from a lack of demand

No matter the form of an initiative — whether a trust, co-op, a collaborative, or something else entirely — it will not succeed if it cannot attract participants. Many data stewardship initiatives face the dual challenge of not only attracting participants, but educating them on how to benefit from (or mitigate their risk towards) data. Beyond that, participants could be deterred by an unclear value proposition, a confusing design, or unreasonable expectations of a participant's available time or expertise.

Story: A confusing consumer data trust

A non-profit builds a data trust for consumer-related data. The trust works on behalf of its members to secure better terms for member data, particularly related to consumer activity and targeted advertising.

Potential users aren't persuaded that the trust is beneficial, don't understand the benefits, or object to their data being sold at all. The trust struggles to find vendors who are interested in purchasing the trust's data, and those that do continue to engage with other, exploitative data brokers. Consumer advocates object to the trust for normalizing the exploitation of information about a person's life and relationships for financial gain.

Story: An opaque privacy app

A non-profit builds an app that claims to help users protect their data from companies. The app consists of dozens of links to optimize privacy settings in the user's social media, email, and other online accounts, and notifications about where the user's personal information has been leaked online. Many users find the setup process to be overly time-consuming, and the benefits of using the app long-term are unclear. The app fails to build a trust relationship with users over time, and most eventually delete it.

Story: Overwhelmed by decisions

An application to give users agency over how their health data — from patient records to wearables data — is developed. For legal and business reasons, the app does not make recommendations on whether users should accept or reject an individual request, or when they should negotiate better terms. Users are quickly overwhelmed by the volume and complexity of potential requests for their data. Some users see a request's presence on the app as evidence that it has been vetted for trustworthiness, and file lawsuits against the app's maker when a deal goes bad. Others begin to reject all requests that come in, or accept only requests that pay, even a nominal amount.

Data stewardship initiatives fail to overcome collective action challenges

Some data stewardship initiatives are built on theories of collective action: given enough users, collectives can negotiate better terms with platforms that handle personal data.

The success of this “data union” approach may be situational and small-scale: where the union has meaningful negotiating leverage, such as monopoly over a data source or a user community, and where the union is only negotiating with a few platforms, instead of many. At scale, however, data unions may share few of the strengths of a labor union, and add unique weaknesses. It is not clear that a "data union" can gain a critical mass of members to successfully negotiate with a platform, especially when large platforms have chosen to leave entire markets with unfavorable legal environments (e.g., Google News in Spain, Uber and Lyft threatening to leave California). Without that critical mass, it's difficult to attract new members.

Perhaps more importantly, at any scale, a data union may not be able to prevent its members from making side deals with platforms they need to use, undermining the union's negotiating power.

Story: Trading body privacy for affordable health insurance

Xander has Type 2 diabetes. He agrees to allow his insurance provider access to fitness tracker and medical device data in exchange for lower premiums, even though the insurance provider refuses to negotiate with Xander's data union. When forced to choose between the data union and a service he needs, Xander leaves the union.

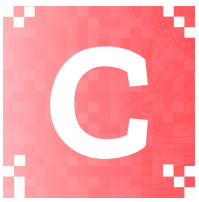

Data stewardship initiatives fail to win user trust

Some data stewardship theories of change rely on trusted expertise: whether in the form of a trusted fiduciary, or in an individual's ability to make informed decisions about how to protect their data.

Initiatives may fail to maintain a trusted relationship with their members. They may engage non-expert fiduciaries, or show bad judgement. Security breaches may further erode an initiative's trust.

Data stewardship initiatives may go one step further: in addition to stewarding and protecting data, they may try to help people analyze data and make decisions about their lives. This risks overreach: a health data steward may not be the ideal health advisor.

To succeed as a movement, data stewardship initiatives need to educate the public on the benefits and risks of data, especially in an age of machine learning. Without that, even successful implementations may fail to counteract a growing public distrust towards data and those who use it.

Story: Data leaks lead to real-world harm

A local government purchases ride-sharing data from a driver-owned co-op. A mistaken response to a FOIA request exposes individual ride data to the public. After a

wave of stalking and domestic violence incidents, the data-sharing program is shut down.

Story: Misinterpreted data

A patient-run data co-op collects sensor readings from fitness trackers and other health wearables to build a crowd-sourced, privacy-protecting fall detection notification system. Although the data co-op is a responsible steward of the data collected, they produce a poor-quality fall detection system. After several publicized incidents of falls that do not trigger the detection system, health officials begin warning people not to use the tool, and users begin abandoning the co-op.

Story: Surveillance without representation

A labor union encourages members to install a monitoring application on their phones, promising to use the data collected to improve their negotiating position with employers. The union invests money and time in maintaining the application and the attendant data, at the cost of grassroots, person-to-person organizing support. Ultimately, much of the data collected by the monitoring application goes unused, harming workers' faith in the union.

Data stewardship initiatives are financially unsustainable

The promises a data stewardship initiative makes are only as good as its long-term financial health. Without good planning, a stewardship initiative's collapse or closure can harm the populations that depend on the initiative in order to run.

Story: Voice data up for grabs after collaboration falls silent

A consortium of local organizations collaborates to build and maintain a voice data set for African languages. The consortium develops a text-to-speech and speech-to-text API that is affordable and has a well-scoped acceptable use policy.

Because the datasets and software used to create the APIs are open and publicly available, competing speech recognition APIs arise that are targeted at the most lucrative use cases. Eventually, the consortium runs out of money to maintain and steward the datasets, and the project is abandoned. Some of the language datasets are adopted by corporate sponsors, who use them to develop proprietary tools. Others are left online for anyone to use, for any purpose. No one is left to enforce the acceptable use policy. A European company uses the data to develop and sell localized voice surveillance tools to governments.

Story: A failed patient hub

A foundation funds a patient advocacy group to build a shared hub for members' health data. After the seed funding runs out, the patient advocacy group cannot raise additional money to develop the hub or pay for its hosting. Eventually, the group's only option is to get funding from a health insurance company. Although the insurer does not explicitly request access to the group's data as a condition of funding, the group's members, many of whom have spent years fighting with their insurers over coverage, become wary of contributing to the hub.

Data stewardship initiatives are unable to resolve conflicts among stakeholders

Like other multi-party initiatives, data collaborations can fall apart when members cannot resolve different or opposing goals, and are unable to negotiate terms for production, use, and control over data derivatives.

Story: Should a neighborhood watch?

A neighborhood data commission fractures over a request to make camera, microphone, and sensor data available to a commercial vendor selling gunfire detection software to the city and its police department. After the request is narrowly voted down, individual business owners and residents install cameras and microphones in their own windows, and attempt to make that data available to the vendor.

Notes and further reading

Arguments in 1A, 1C, and 1E are inspired in part by Catherine Leviten-Reid and Brett Fairbain, "Multi-stakeholder Governance in Cooperative Organizations: Toward a New Framework for Research?" Canadian Journal of Nonprofit and Social Economy Research, Vol. 2, No. 2, 25-36 (Fall 2011).

Leviten-Reid and Fairbain review available empirical evidence on multi-stakeholder cooperatives, which they argue suggests some indication of successful governance. Citing Ostrom, they highlight four processes that contribute to successful governance: 1) that all actors are involved or have the opportunity to be involved in rule-making; 2) that individuals have a clear understanding of the system they are a part of and the rights and obligations they have; 3) that conflicts are resolved quickly, due to in-built mechanisms for identifying and addressing them; and 4) that members perceive a fair relationship between what they invest and the extent to which they benefit.

The examples in 1A about users being overwhelmed are inspired in part by the WEF's report "The Internet of Bodies is here: tackling new challenges of technology governance." More generally, the report explores how body-connected sensors and devices raise new governance challenges. The privacy app example in 1A is also inspired in part by Kate Cox's article in Ars Technica, "Unredacted suit shows Google's own engineers confused by privacy settings."

The argument in 1B is inspired in part by the work of Mancur Olson and others, who discuss the challenges of organizing large groups of people with diffuse interests, especially when pitted against smaller groups with more acutely defined interests. See

also Amy Kapczynski, *The Access to Knowledge Mobilization and the New Politics of Intellectual Property*, 117 Yale Law Journal 804, 811 (2008).

The fall-detection example in 1C is adapted from M. Schukat, et al., "Unintended Consequences of Wearable Sensor Use in Healthcare," Yearbook of Medical Informatics (2016).

The examples in 1D are inspired in part by the RadioShack bankruptcy case, where the company attempted to sell consumer data despite promising that it would never do so. Intervention from the FTC and 38 states ended up limiting the extent of the deal.

The example in 1E is partially a riff on controversies related to Ring, NextDoor, and ShotSpotter.

Part II: Exclusion

Data and technology projects continue to exclude, discriminate, and marginalize. New data stewardship models create new barriers to meaningful participation, or reinforce old barriers.

Data stewardship models may fail to lead to a more equitable digital world.

For one, data stewardship models do not guarantee equitable governance or operations. Without constant effort and attention, data stewardship initiatives may exclude and marginalize, and may fail to protect vulnerable community members.

For another, data stewardship models may be unable to prevent the weaponization of data and inferences against vulnerable populations. In many cases, the best way to keep data safe is still not to collect it at all.

The stories in this part highlight how data can be used to exclude, discriminate, and harm people. While not all of the stories here are caused by data stewardship initiatives failing, an initiative fails when it allows stories like this to form.

Data stewardship initiatives may be governed inequitably

Data stewardship initiatives that are built around community governance run the risk of inequitable governance. Community members may not have equal ability to participate in the governance of a co-op or collaborative. They may lack time, training, or opportunity to make informed decisions about what to do with their data. Stewardship initiatives have an obligation to lift members up: to provide training, to provide opportunities for input into decision-making, to provide protection for vulnerable members. Initiatives that fail to meet those obligations will only exacerbate inequality.

Story: Barriers to participation and governance

Parents of children with chronic conditions form a data co-op to help safeguard their children's health data. Membership in the group is primarily via word-of-mouth. The group primarily holds meetings during the workday, which makes it more difficult for parents who work to attend the meetings and meaningfully participate in the group. As a result, the co-op's patient population is richer and whiter than the national patient population. After the co-op contracts with a technology company to produce an algorithm to triage and predict severe disease course, audits reveal that it produces worse clinical outcomes for Black and Hispanic patients.

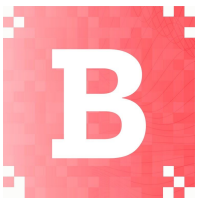

Data stewardship initiatives may further data protection inequities

Where data stewardship initiatives attempt to negotiate directly with platforms — for better terms, for new features — they risk creating new inequalities. Even now,

platforms discriminate based on national and local policy differences: GDPR's rights and protections are not always extended to non-European users.

Data stewardship initiatives that negotiate on behalf of communities may further this patchwork. Assuming platforms engage at all, a wealthy group of users in the United States may receive better data protection terms than, say, civil society groups representing rural minorities in India.

What does this mean? The success of some data stewardship initiatives may depend more on the policy and economic context it lives in, rather than the stewardship model itself. This may confound efforts to scale stewardship models beyond the local level.

Story: Data rights arbitrage

Rideshare drivers in the UK successfully use subject access requests to build a detailed database of wage data. In response, rideshare companies negotiate a more equitable wage model that minimizes discrimination between individual drivers for the same ride. The driver collective and the rideshare companies sign a non-disclosure agreement.

This model is deployed only in the UK. In countries where individuals do not have the right to make subject access requests (such as the US), or in countries where drivers have less economic leverage (such as Thailand), the rideshare company continues to use its discriminatory wage model.

Data stewardship fails to prevent algorithmic discrimination

Data is potential. A single dataset or a single inference or algorithm may have multiple potential uses: from diagnostic to discriminatory. It is increasingly difficult to control these uses and reuses, even with better data stewardship. Achieving accountability for data-related discrimination may require more wholesale policy changes than a stewardship initiative can manage on its own.

Story: Voice-powered discrimination

By analyzing voice records of interview data from a longitudinal study, researchers develop an algorithm to detect early signs of cognitive decline based on biomarkers in voice data. The algorithm is developed with full patient consent, and effectively preserves the privacy of the initial study population.

After a blood-based testing protocol proves to be more reliable for clinical diagnosis, the algorithm is openly licensed. Soon, the algorithm is deployed in automated job screening tools (to filter high-risk candidates), insurance support hotlines (to re-price insurance for customers), and video-based social media (to drive targeted advertising).

Data stewardship models fail to prevent the weaponization of data against vulnerable groups

Data stewardship initiatives are still vulnerable to power asymmetries. Even if a data steward succeeds in protecting members from commercial exploitation, most, if not all, are still vulnerable to legal exploitation: from government seizure to legal action. Here, data stewardship is powerless to stop the weaponization of data against vulnerable groups, whether in Hong Kong, the United States, or anywhere in between. Often, the safest option is not to collect data in the first place.

Story: Weaponizing a storytelling program.

With the help of volunteers, an immigration activism group begins to compile a storytelling corpus. To encourage candor, the group promises to embargo stories from release until after the storyteller has died. To keep their promise, the group holds the stories in trust with a third-party trustee.

Five years into the program, a new administration comes to power that takes a harder line on immigration. Immigration enforcement successfully obtains a warrant for the storytelling data, and begins to reidentify the participants. Several are arrested.

Notes and further reading

The story in 2A is inspired by a racially-biased triage algorithm deployed across American hospitals.

The argument in 2B is inspired by several pieces on regulatory arbitrage. See, e.g., Ryan Calo and Alex Rosenblat, "The Taking Economy: Uber, Information, and Power," 117 Columbia Law Review 6; Brishen Rogers, "The Social Cost of Uber," 82 University of Chicago Law Review Online (2017). See also research from Itzhak Ben-David, Stefanie Kleimeier, and Michael Viehs on how companies take advantage of different national pollution regulations.

The argument in 2C is inspired in part by Nathaniel Raymond's piece "Safeguards for human subjects research can't cope with big data." (Nature, 2019) The story in 2C is inspired in part by actual research on using voice data to detect cognitive decline. For more on data-driven price discrimination, see, e.g., Silvia Merler, "Big data and first-degree price discrimination," Bruegel (2017); Christopher Townley, Eric Morrison, and Karen Yeung, "Big Data and Personalised Price Discrimination in EU Competition Law."

The argument in 2D draws from the work of Yeshimabeit Milner, ("We will not allow the weaponization of COVID-19 data"), Mutale Nkonde ("Congress Must Act on Regulating Deepfakes"), Joy Buolamwini ("We must fight surveillance to protect Black lives"), and Virginia Eubanks. The example in 2D is inspired by the Boston College IRA Tapes Project.

Part III: Entrenchment

Power asymmetries persist and grow worse. Entrenched interests co-opt, outcompete, or crush alternative models.

There is a potential mismatch between the goals Mozilla articulated in commissioning this brief and its subsequent analysis of data stewardship. While Mozilla seeks to shift economic and social power, many of the data stewardship initiatives the landscape analysis cites have more modest goals: to make data available, to improve privacy, to afford individuals more situational agency over how data is used.

This serves to underscore that data stewardship is an insufficient frame to truly shift power. Worse, if data stewardship is deployed as an antidote to platform dominance and surveillance economies, it threatens to validate the exploitation of human lives and relationships — the true source of data — for financial gain or "value." The idea that we can conduct this exploitation "responsibly" is fiction, and avoids the more important question of whether we should do it at all. There is nothing about search, social networks, or our online lives that requires our consent to data-driven exploitation. It is merely the price that has been set for us.

At a minimum, to shift power around data would require uprooting the targeted advertising model that much of big tech relies on, either directly or indirectly. Moreover, this would indirectly serve as an existential threat to Mozilla's current funding model, which relies heavily on commercial search engines.

This part reviews failure models that relate to entrenchment: how existing power asymmetries can co-op, resist, and crush data stewardship initiatives.

Entrenched powers refuse to engage with data stewardship initiatives.

Platforms and governments may choose not to engage with data stewardship initiatives, or attempt to undermine their efforts to force change. Data availability alone may not be enough to force accountability where none exists.

Story: Rideshare company plays hardball

Rideshare drivers in the UK use subject access requests (SARs) to build an alternative database of driver wage data, and hope to use that to negotiate better terms with a rideshare company on behalf of all drivers.

Borrowing from litigation and insurance company playbooks, the rideshare company begins taking measures to increase the effective cost of the SARs. Requests are routinely delayed or lost. Data is delivered via PDF or mail courier. Data formats change at random. Location data and ride data are hashed and deidentified, limiting the data eligible for disclosure.

In the meantime, the world has changed. Ridership since the pandemic has cratered. Some of that drop turns out to be permanent, as people work from home, go out less, and have less money to spend, while many of the businesses people would take a rideshare to — restaurants, bars, entertainment venues, clubs, retail — have closed. Many drivers start looking for new work, shrinking the pool of potential participants. Moreover, the data from pre-pandemic are less useful, because they reflect an entirely different set of market conditions and social habits. With its financial losses mounting, the rideshare decides to leave the UK rather than negotiate with drivers.

Story: City government ignores citizens and their data

A group of neighborhood commissions begin to adopt a software tool for citizens to report potholes, broken streetlights, and other non-emergency maintenance issues. They hope to pressure the city into taking a more proactive approach to maintenance in

their neighborhoods. The city government ignores the software. Six months and one unsuccessful lawsuit later, more than 90% of the citizen-reported issues are still unfixed. Use of the tool begins to decline.

Data stewardship models are co-opted by technology platforms

Many data stewardship models are new, untested, and relatively unknown. This leaves them vulnerable to co-opting by technology companies or other hostile actors. Even within a “successful” data stewardship initiative, big tech companies may seize a disproportionate share of the economic benefit.

Story: An urban data trust free-for-all

A technology company sponsors a data trust for a smart city. The default rules of the data trust make city sensor data publicly available, including from the technology company's competitors. The trust does not say whether software developed from those datasets must remain open. Because of its size, the technology company is better positioned to develop proprietary software based on the open data, and sells it back to the city.

Story: Profiting from school data analytics

A company manages an analytics platform and data collaborative for a dozen school districts across the American Southeast. The collaborative is designed to allow school districts and their partners to better understand long-term student outcomes. While most access to student data is governed by the collaborative's IRB, a loophole in the agreement allows the managing company to analyze the student data independently without going through the IRB, develop algorithms based on their analysis, and retain

all of the intellectual property related to those algorithms. The company later sells itself for $350 million. The school districts get nothing.

Story: Infiltrating a driver co-op

A rideshare driver data co-operative allows anyone who has driven a minimal amount of rides to join. The co-op allows members to vote on major decisions. A bloc of rideshare company employees pose as drivers to join the co-op, and begin pushing company-friendly policies.

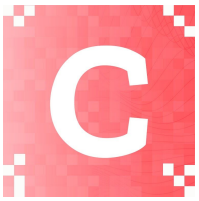

Abuse or poor judgment by fiduciaries and trusted parties is difficult to identify and remediate

Data stewardship models that entrust control to a fiduciary or other trusted party are vulnerable to abuse or poor judgment. If conducted on a platform, this abuse may be difficult to identify. Legal uncertainty about how fiduciary duties are applied to data may make it difficult to distinguish between negligence and (allowable) bad judgment.

This highlights a broader challenge: trust doesn't scale. What works in a smaller community or network may not scale up to a national or global level.

Story: Intermediary fraud in benefits system

A national government develops a system to help citizens manage where their personal data is used. The system is integrated with a national identification and benefits system, and requires a mobile phone to use. People who have low-literacy or poor access to mobile phones must rely on intermediaries to access services on the system. An intermediary could be a (usually male) head of household, a neighbor, or even a complete stranger. Intermediary fraud becomes common, and difficult to uproot:

because most information about the system is distributed via phone, victims are either unaware or unable to report issues or seek accountability.

Story: Contact-tracing workarounds

A privacy-preserving contact tracing app begins to see wide use in a country. Fearful of infections, businesses begin to require citizens to show app-based confirmation of "no exposure" before they are allowed to shop. A secondary market for phones and "clean" apps becomes rampant, undermining trust in the application.

The surveillance economy cannot be normalized

Data stewardship initiatives are unlikely to reduce the large-scale exploitation of data. In part, this is because data is a nonrival good: one person's use of data does not prevent anyone else's use of it. This means that unlike natural resources or other physical goods, data stewardships may not be able to maintain a monopoly on data they protect. The practice of building "data doubles" — a model of a consumer based on indirect inferences about them — may continue unabated.

By attempting to build "guardrails" on the surveillance economy, data stewardship initiatives may actively make things worse. Allowing individuals to buy and sell access to their data warps privacy from a right into a commodity, and may enable new forms of exploitation.

We have succeeded in building an economy where it is normal to buy and sell information about how a person feels, about who they love and hate, about where they go, and what they hope and fear. Building a just world will require far more than haggling over the price.

Story: Data doubles, subprime data

A data co-op for individual consumer data is disrupted because it cannot maintain a monopoly on user data. Whether through phone applications, internet browsing history, subscription services, or financial transactions, data brokers have more than enough sources of information to target and market to users, and most decline to purchase access to individual data.

Ultimately, the only users who make money from the co-op are those who sell extremely sensitive data that is difficult to get otherwise, such as emotional and health data. These “subprime” data purchases are targeted towards people with low income and poor credit. In some cases, personal data is used as collateral against a payday loan. Ultimately, data is re-sold to tenants, employers, businesses, and governments.

Notes and further reading

The arguments in the introduction to this section are drawn from, among others: Nick Couldry and Ulises A. Mejias's "The Costs of Connection;" Lina Khan and David Pozen's paper on information fiduciaries and the related symposium on the Law and Political Economy Project (in particular Julie Cohen's contribution); and Ruha Benjamin's work on "Informed Refusal."

The urban data trust story in 3B is inspired by one of Sidewalk Labs' proposed "trusts" for Waterfront Toronto.

The school data analytics story in 3B is very loosely inspired by the sale of Flatiron Health.

The intermediary fraud story in 3C is inspired by Anita Gurumurthy, Deepti Bharthur, and Nandini Chami's research on exclusionary practices related to the Jandhan Aadhaar Mobile (JAM) payment platform in Rajasthan, India.

The data doubles argument and story in 3D is inspired by Kevin D. Haggerty and Richard V. Ericson's work, "The surveillant assemblage."

Afterword

Often as not, data are stand-ins for people: for their desires, their anxieties, their livelihoods, their secrets. This has made it easy to frame the exploitation of people as a benign product feature: “we will identify depressed people and target them with advertising” becomes “we will analyze social media data for emotional affect.”

Data stewardship initiatives risk falling into a similar trap, as they attempt to productize trust, good governance, and stewardship. Initiatives may use the language of data stewardship for branding purposes (see, e.g., "data trust"), make promises that are easily broken or dodged ("we'll never sell your data"), or overstate their replicability and scalability.

In truth, nothing scales like exploitation. It's possible — even likely — that none of the data stewardship initiatives described in Mozilla’s research will scale, and that few success stories will ever be replicable. This isn't a sign that the initiatives are flukes, or that their underlying models are defective. Rather, human relationships — trust between people — can't be copied from one initiative to the next. And a good data stewardship model can only take an initiative so far. Regardless of their form, data initiatives must make difficult choices about whether to collect and keep data; about who can access data; and about how to use and analyze data. They must do the hard work to continuously earn and re-earn people's trust, as technology advances and grows more difficult to understand. And even then, it may not be enough to stave off failure, or prevent harm.

Asking how to shift power through data governance, then, is not enough. The challenge before us is much bigger: to remake our relationships with one another, with the digital communities we join, and the leaders we trust. To successfully meet the moment requires far more than innovative models for managing data. Truly shifting power requires — *demands* — a policy and popular movement that rebuilds public understanding of data, rehabilitates digital communities, and redefines our digital economy from the ground up.