

**BY NATHAN KINCH AND MIKE HALEY ON OCTOBER 14, 2020**

Following on from [Predictive Tech and Data Ethics: Part 1 – the Pros and Cons](#) our intention with this post is to help individual practitioners and business leaders better understand our rationale behind why data ethics matters and how it can be done better. This isn't about feel-good principles-based statements, it's about effectively operationalising data ethics. Doing this, we hypothesise, has the potential to increase data and brand trust. We do this by describing a step by step implementation process. We support this with concise commentary regarding Autodesk's direct experience with this approach.

Although [data ethics](#) is a relatively new focus for most organisations, "formal" initiatives appear to have had little positive impact on consumer trust. In fact, Pew Research Centre's [most recent privacy study](#) shows that US consumers are more concerned than ever about how public and private organisations are using data about them. Study respondents believe their data is less secure, that data collection poses more risks than benefits, and that it is not possible to navigate their daily lives without being surveilled.

This consumer discomfort has led to a [data trust deficit](#), so that the trust an individual places in a brand and the trust they place in its data practices are different. In most cases, data trust is lower than brand trust.

Although this is a complex issue, we believe that the limited scope of the common principles-based approaches to data ethics – which are primarily limited to the ethics of data collection and use – make it difficult for organisations to improve the trustworthiness of their practices. It's a relatively narrow approach that shifts focus away from the bigger issues that

have an impact on trust, like incentives misalignment. Perhaps this is why many are calling out Big Tech for [ethics washing](#).

We've, therefore, extended the scope of data ethics initiatives to focus on an organisation's entire operating model, considering the ways in which data has an impact on decision making, specific actions and real-life outcomes. This includes everything from an organisation's business model and incentives structures, through to product design and development, marketing and customer support and even employee training.

In addition to this extension of scope, we propose that organisations should move beyond feel-good statements and commit to operationalising data ethics frameworks. A data ethics framework, as we define it, is the consistent process an organisation executes to decide, document and verify that its data processing activities – both the intent and real-life outcomes of those activities – are socially preferable.

## From Principles-based Approaches to Operational Frameworks

Shifting focus away from ambiguous principles towards more strategic, operational and accountable frameworks requires leadership courage, meaningful investment, and time. This is a significant challenge, particularly in large, complex organisations.

However, the approach we advocate can be approached pragmatically. Here's how.

### Step 1: Build Your Core Team

From experience, we know there are many people within an organisation who care deeply about the broad and nuanced topic of data ethics. So, before an organisation formally commits to and funds a data ethics framework, plenty of grassroots work can be done.

The foundation of this is establishing a diverse, influential, and motivated group of team members to act as early evangelists. These individuals must be close to the problem, so consider representatives from product management and marketing, research and development, sales and business development, operations and people and culture.

Identify these individuals. Communicate the intent. Get buy in and start making progress.

It's worth noting that a core assumption here is that an individual or small group of individuals kicks off this process.

## Step 2: Define Your Purpose, Values and Principles

Most organisations have an explicitly documented purpose. Many have documented values and principles. If this is the case, the core team can build upon this to establish foundational content that will define what is good (values) and what is right (principles). This may involve internal workshops and dedicated work streams to coordinate collaboration that produces specific documented output.

The Ethics Centre, whom we've worked with on past data ethics projects, [has a framework for this](#). If you've yet to define values and principles, this is a great model to consider using.

## Step 3: Raise Funding at the Right Time

Without appropriate resourcing and capital allocation, the detailed work required to design and progressively test a data ethics framework cannot be done.

With your core team, draw on primary and secondary data to build a succinct and compelling business case for the strategic importance of this work. It should include both qualitative and quantitative data sources, and, where possible, also reflect the organisation's documented strategy. The more closely your proposed initiative aligns to existing strategic intent, the greater likelihood the initiative will receive the early funding it requires.

#### Step 4: Conduct Internal Research

The success of a data ethics implementation rests on changing behaviour. Therefore, internal research needs to focus on current attitudes, behaviours and related incentive structures. We use three interrelated approaches to support stakeholder interviews:

1. Importance vs Satisfaction: Using a Likert Scale (1 – 7), we conduct qualitative surveys to understand the importance that key stakeholder groups assign to this initiative, as well as the satisfaction they assign to the way the organisation currently handles data ethics challenges and opportunities.
2. B=MAP: [The Fogg Behaviour Model](#) is used to develop an understanding of the intersection between motivation and ability as axioms that influence behaviour.
3. The Switching Canvas: An adaptation of [ReWire Group's switching formula](#) is used to identify and influence the forces that encourage and inhibit specific switching behaviours.

These approaches enable us to identify the likely needed behavioural changes if the organisation is to act on its data ethics intentions.

### Step 5: Define the Desired Future State

This, although highly speculative, is an important activity. It helps the organisation to begin to define ways in which data ethics has created a positive impact. The more explicitly this future state is defined – what the organisation does and does not do, the impact this has had on operations, performance, culture and stakeholder experience – the better.

This end state may never be achieved, that isn't the point. The point is to set an ambitious target, something that inspires and further motivates those involved.

### Step 6: Define the Change Pathway Towards the Desired State

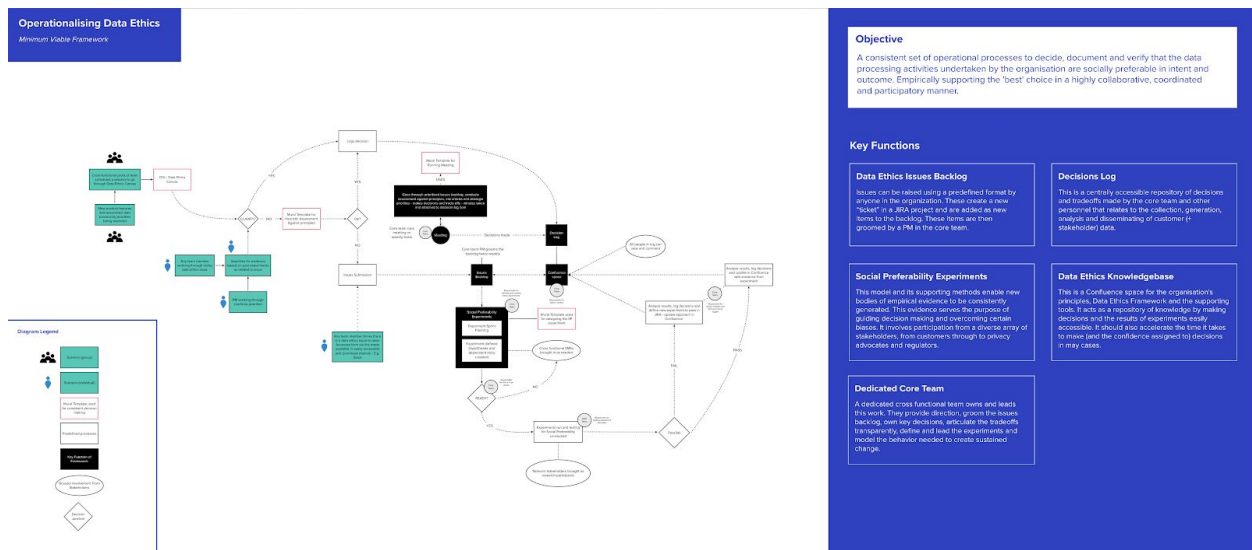
With an understanding of the current behaviours, workflows and practices, and a clearly articulated longer term vision, you can begin to hypothesise the actions that might close the gap between a current and desired future state.

This becomes the foundation of your strategy or plan. It defines the focus areas, specific tactics, measures of progress and incentive structures that will consistently support the desired behaviours and outcomes.

### Step 7: Design the Hypothetical Framework

In order to operationalise the purpose, values and principles, and thus the desired behaviours you've identified, you need a consistent way to make decisions, document those decisions, and verify that they are socially preferable in intent and outcome. This is what your data ethics framework is designed to help you achieve.

This is documented diagrammatically so that the idea flow, decision-making processes, feedback loops and related functions can be better understood. This visual representation of the framework can be supported by additional detail that describes specific framework functions, like the way backlog items are created, groomed and prioritised, or the way social preferability experiments are conducted. Commonly used formats like Business Process Modelling Language can be used if preferred.



Your data ethics framework should be documented diagrammatically

## Step 8: Build, Groom and Prioritise the Backlog

The data ethics issues backlog is a "whole of organisation" process, designed to be accessible to everyone.

Note that there's a crucial relationship between the backlog and knowledge base. Team members may identify an issue or opportunity that they don't think has already been identified and acted upon. Part of the backlog creation process is searching the knowledge base: if nothing exists, a new issue can be created.

Grooming and prioritisation also have a clear and consistent process. This is typically the responsibility of core data ethics team members.

## Step 9: Conduct Your First Social Preferability Experiment

Our approach to data ethics assumes that ethical decisions are more effective when the individuals, groups, communities and organisations affected by them are actively involved in the decision-making process. This is why social preferability is the key focus of the frameworks we design. This approach augments principles-based decision making with empirical data from experiments and actively and consistently engages an organisation's key stakeholder groups in the decision-making process.

Here's a basic overview of how to run an early-stage social preferability experiment.

Let's start with a use case. You want to deploy a new model intent on improving customer service requests. At the time of designing this experiment and proposing this new activity, you don't have explicit permission from your customers to use data for this data processing purpose, it's a new activity. You're proposing it because you believe it will deliver direct value to your customers and create operational efficiencies for your business.

There are two key areas of focus for the experiment:

1. This first is that you need customers to give their explicit permission to use the data from their support requests to train new models.
2. The second is that you need to show the outcome this new data processing might enable for them. You need to communicate and demonstrate the new value of the proposed activity.

To do this you might develop a basic prototype of the permission flow, clearly articulating the why, what, and how of this new data processing activity. You would also build the desired outcome into the prototype. This would enable research participants to better understand the impact of your proposal. You do this because social preferability testing seeks to understand the support key stakeholders have for both the intent and outcomes of data processing activities.

Once you have the prototype, you design, recruit for, and run a tight research program. We refer to this as outcome-focused usability paired with contextual inquiry. It's a hybrid research method that helps to gather proxy quantitative data by simulating real-life usage. These metrics are then supported by qualitative and attitudinal data as a result of the researcher-led contextual inquiry.

The simplest approach to this is to embed [Likert scale questions](#) at different stages of the research participants' experience, like directly after someone takes an action to grant permission. The framing of the prompting question is important. If you're engaging other stakeholder groups like internal team members from legal or compliance, it's important they approach this research session from a customer's perspective. The same would be needed if you were to engage regulators in this research. We expect these stakeholders to contribute their professional perspective, but it's more important they empathise with the individuals and groups the proposition seeks to serve. Additional qualitative data from the sessions can be used in analysis stages to put the self-asserted scores in context.

In this scenario you're measuring a single dimension where 1 = socially unacceptable, 4 = socially acceptable and 7 = socially preferable. If you want to improve the rigour of this approach you might consider exploring the use of social value orientations. However, for the



purpose of building some foundational evidence to support ethical decision making, we have found this approach the most time and cost effective.

## Step 10: Socialise What You've Learned

[Progress](#) is an important motivator. Celebrating small wins can go a long way.

Whenever you conduct an experiment, add the process, artefacts, and outcomes to the data ethics knowledge base. We also advocate that it's communicated directly to core and secondary data ethics team members and influencers, and that the results are broadly communicated via a channel that's most likely to encourage active engagement, such as email, or a messaging tool like Slack.

This demonstrates that data ethics is both practical and impactful, through helping establish precedent that can support faster and more consistent ethical decision making in the future.

As an addition, these details could be shared openly with key stakeholders like shareholders, customers and regulators to support a proactive approach to openness and accountability. An example of this can be seen in [Greater Than Learning's Key Decision Log](#).

## Step 11: Refine the Implementation Plan

It's likely that your implementation plan has been in progress for some time. It can now be refined with greater confidence as a result of experiments conducted and lessons learned.

We suggest designing the implementation plan like a product roadmap. This includes dual-track discovery and delivery workstreams and confidence interval based detail, meaning actions in the next few weeks are explicitly defined and actions taking place beyond that are more loosely

defined. This takes into account the reality of business planning and enables team members to focus on the defined activities most likely to be acted upon.

## Step 12: Progressively Implement

Data ethics is never done. That's why thinking about your data ethics framework as a platform or product is useful. It should be designed to offer extensibility, while also being fit for the purpose of your current context.

Over time – based on what you test, learn and share – the platform that is your data ethics framework will evolve. You will become more adept at consistently exhibiting the desired behaviours. The actions the framework supports will become more embedded into strategic and tactical workflows.

Although this may now seem like a fairly straightforward point, operationalising data ethics isn't about a policy or even about your principles. It's about the behaviours the people within the organisation exhibit on a daily basis. It's about changing and adapting your behaviours of today to support the transition towards your desired future state.

## In the Real World

Autodesk has realised that earning the trust of its customers begins with a technology and business strategy rooted in data ethics.

The company balanced grassroots action with a highly coordinated strategic approach at the executive level. The grassroots effort involved selecting key proponents from across the company to drive specific concrete data ethics successes. These successes help establish data ethics' reputation as a business-value driver as opposed to a regulatory framework. Meanwhile,

the executives identified a long-term business goal that would benefit from a data ethics framework. This created a top-line driver that, when combined with the concrete results being produced by the initial applications of the data ethics framework, is steadily transforming Autodesk into a major adopter of data ethics, both as a customer value proposition and as a cultural value.

## Conclusion

It's our belief that the most trustworthy organisations will develop the most meaningful relationships with their key stakeholders. The quality of these relationships will likely lead to more effective, trust-based data sharing, the development of better products and services and more sustainable business practices. A thoughtfully designed and pragmatically implemented data ethics framework has the potential to make this a reality.

## Read Part 1: Predictive Tech and Data Ethics: the Pros and Cons

In Part 1, we considered how product managers can begin to address the risks and ethical issues associated with predictive technology. In it we covered:

- The importance of trust
- Understanding the issues with data
- Developing an ethical product practice
- Simple tips on how to get started

Missed it? [Read Predictive Tech and Data Ethics: the Pros and Cons now.](#)



## About

### Nathan Kinch

Co-founder of [Greater Than Learning](#) and CEO of [Greater Than X](#) Nathan is a 3x entrepreneur, speaker and writer, impact-focused startup investor, wannabe athlete and plant-based home chef. He's spent most of his career designing organisations for the qualities of trustworthiness. He's worked with federal governments, major corporations, fast-growing startups and research and policy institutes around the globe. He lives on the sunny Gold Coast (Australia) with his life partner and best friend, crazy cool 2-year-old daughter and their tiny four-legged friend, Nala.



## About

### Mike Haley

Mike leads a world-class team of researchers, engineers and other specialists, at [Autodesk](#), who are exploring the future of how people will design and make things. His team works in the areas of: Machine Learning, Human-Computer Interaction, Simulation and Systems, Optimization, Geometry, Visualization and Robotic Control and Sensing. These areas of science are brought together by Industry Futures teams, who work closely with customers and partners, to develop novel solutions to problems in Design, Engineering, Manufacturing, Architecture, Construction and Entertainment.



## About

### Mat Mytko

Co-founder of [Greater Than Learning](#) and Chief Platform Officer of [Greater Than X](#) Mat is a humanist and generalist, father, husband, amateur gardener and extreme sports enthusiast. A builder at heart with an engineering mind and a penchant for creative expression. He's spent the last half of his 26-year career working in and building startups in product and platform roles. He has project experience in diverse fields such as digital identity, vocational

education, e-learning systems, behavioural design, decentralised technologies, system design and organisational transformation. He resides in Sydney with his wife, 2 children, 4 dogs and loving plants.