TAKSHASHILA
INSTITUTION

DISCUSSION DOCUMENT
JULY, 2017

BEYOND CONSENT: A NEW PARADIGM FOR DATA PROTECTION

DISCUSSION DOCUMENT 2017-03

July 2017

By RAHUL MATTHAN[1]
Fellow, Technology & Policy Programme
The Takshashila Institution

The Takshashila Institution,
Bangalore, India

This discussion document can be cited as "Rahul Matthan, *Beyond Consent – A New Paradigm for Data Protection,* Takshashila Discussion Document, 2017-03".

EXECUTIVE SUMMARY

Data protection is confined by consent (**Consent Model**). Once a data subject's consent is obtained, a data controller is free to collect, process and use such data for the specified purpose and will not be liable for any consequences that might result from its actions. This places the onus on an individual to be aware of the terms of data access to which he is providing his consent to. This clearly benefits data controllers more than data subjects.

This is inadequate in the interconnected, data-reliant world of today. Given that India will be working on a formal law on data protection in the near future, it is imperative that it relies on an alternative to the Consent Model in order to protect the interests of data subjects.

We believe that a rights based model (**Rights Model**) will help secure the interests of a data subject sharing his data with data controllers. This Rights Model assures to every individual, an inalienable right over his personal data. Any data collector that wishes to access a data subject's personal data must ensure that they do so in a manner that does not violate this inherent data right. This Discussion Document sets out the contours of such a rights based model (Rights Model) as a substitute for the Consent Model. The Rights Model has the following features:

- It **assures a set of data rights** that are available to everyone.
- It shifts the burden of evaluating the privacy risk to personal data away from the data subject and onto the data controller, **forcing the data controller to be mindful** of its processes for data collection, processing, transfer and storage. The Model **applies equally to the State** when it collects / processes personal data.
- It **focusses on the harm caused to the data subject due a violation** of his data rights, and offers a remedy to him regardless of whether or not he has consented to the terms of a privacy policy. Once a harm is proved, the responsible data controller will be liable for the harm caused to the data subject.

SHORTCOMINGS OF THE CONSENT MODEL

There are three significant reasons why the Consent Model is no longer feasible:

1. The Consent Model is inadequate and causes "consent fatigue"

The Consent Model sufficed earlier because there were limited reasons to collect data and few alternative uses to which it could be put. Once collected, data was static and rarely transferred out of the organisation. Thus, it was easy for data subjects to know exactly what data was being collected and to what purposes it would be put, enabling them to make informed choices. In this context, the Consent Model was both feasible and adequate.

This is no longer the case. Today, data is collected, processed, transferred and consumed in far too many ways to comprehensively enumerate. Our online activity is logged, shopping preferences recorded, recruiters can access our employment histories as well as our social media activity. Every financial transaction we undertake is tracked and correlated against location, age and time of day, offering insights into our personality that even we are unaware of. We are surrounded by smart devices equipped with sensors and cloud intelligence that track our activities and log what we do and how we do it.[2]

We consent to this extensive data collection by signing standard form contracts that are dense and complex, making it difficult to effectively assess the implications of agreeing to the terms set out therein.[3] This, combined with the sheer number of contracts we end up signing, leads to consent fatigue. This results in diminished consent, where we agree to the terms of service and provide our consent without bothering to read the details of what we are consenting to.[4] According to a paper published in 2008, if everyone actually took the time to read privacy policies thoroughly, the national opportunity cost of the time spent on reading privacy policies in the USA would, at that time, have exceeded $781 billion.[5]

2. The rise of interconnected databases increases the need to protect individual data

Modern databases are designed to be interoperable, allowing them to interface with other datasets through APIs. This interconnectivity allows data controllers to layer multiple sets of data and generate powerful insights about data subjects. Privacy policies are now routinely modified to include consent to such interchanges of personal data.

As discussed earlier, it is very difficult to assess the impact of a privacy policy in the context of indirect data collection. Evaluating the impact of interconnected datasets, where the insights gleaned are often unpredictable, is virtually impossible.

3. The rapid increase in data transformation threatens personal privacy

In India, there is no legal requirement to obtain prior consent when collecting non-personal data. Machine learning algorithms are now capable of making connections between seemingly innocuous fragments of data, spotting patterns and building accurate personal profiles. This effectively converts disparate data elements into sensitive personal information. In these circumstances, consent would be insufficient for safeguarding ourselves against the harm that can result from the use of deep learning algorithms.[6]

It bears mentioning that machine learning algorithms and neural networks are designed to function without human involvement.[7]. Consequently, even the most skilled data scientists will not be able to predict how these algorithms process the data provided to them.

FUNDAMENTAL PRINCIPLES BEHIND THE RIGHTS MODEL

To ensure that data subjects are not denied the rights over their data, and that the data controllers are held to account for any harm to privacy, the Rights Model should be founded on the principles of accountability, autonomy and security.

Any viable alternative to the Consent Model must address the issues outlined above. It must ensure that data controllers, who have access to personal information of data subjects, remain accountable for the harm they cause regardless of whether they have obtained consent from the data subject. The model must also ensure that, since privacy is a personal boundary, each individual should have the autonomy to determine his own boundaries when it comes to privacy and in doing so, have the ability to circumscribe the uses to which data controllers can put his data.

With this in mind, the proposed Rights Model must be guided by the following basic principles:

1. Accountability

Data controllers must be responsible for the data they have under their control. If a data subject comes to harm as a result of a security breach or from the manner in which the data controller processed the data, the latter must be held accountable. This liability must exist regardless of any consent given by the data subject.

Importance of the Accountability Principle

The accountability principle assigns a higher degree of responsibility to a data controller than the Consent Model. In doing so, it addresses the shortcomings associated with principles that often find mention in the privacy statutes of other countries. These principles are often extensions of the notion of consent and include:

- The idea of *choice*, i.e., the notion that a data subject should have the choice to determine whether or not his personal data should be collected and processed. This presupposes that data subjects have the ability to understand the implications of allowing their data to be collected and processed.
- The concepts of *purpose limitation* and *use limitation*, i.e., data collectors must ensure that data is collected for a specific purpose and once collected, processed to the extent that fits the purpose. Since it is virtually impossible to obtain meaningful informed consent, limiting use and purpose in reference to what the data subject has consented to is meaningless.
- A *restriction on the disclosure* of personal data to other people or entities without consent. This restriction derives from the notion that the way to protect personal privacy is to limit the dissemination of data unless consent has been obtained from the data subject. This notion assumes that data subjects are able to assess

whether or not a given disclosure or transfer can have an impact on their personal privacy.

2. **Autonomy**

All data subjects should be able to exercise autonomy over their data. Since, in light of modern technologies, it is not possible to effectively prevent the collection of data, every data subject should have the ability to limit or restrict the manner in which data once collected is processed.

3. **Security**

Data must be treated securely at all times, from its collection to ultimate processing and use. Even if no harm is caused due to the loss of data as a consequence of a security breach, the data controller must be punished for non-compliance with fine or imprisonment, as appropriate.

IMPLICATIONS OF THE RIGHTS MODEL

Having articulated the broad principles on which the new data protection law should be based, we will now spell out in a little more detail the contours of this proposed legislation.

1. **Every individual has certain inalienable rights over their personal data**

Data subjects will have a set of data rights. A data controller will be obliged to ensure that these rights are not violated when it deals with a subject's data. While there will be no restriction on the collection or processing of data pertaining to a data subject and consequently no need to first obtain the consent of that person before collection, the new framework will be based on holding the data collector accountable for any harm that it causes.

Thus, the data collector will have a fiduciary responsibility in respect of the data under its control and will be liable for any consequential harm caused to the data subject.

We propose that personal privacy should be protected through the enactment of a statute that guarantees a set of data rights to all persons. These rights, encompassing both personal and non-personal data, can be exercised by the data subjects against anyone who processes or controls their information. Unlike the previous framework, where data subjects' rights were available only with respect to persons they had contractual relationships with, the new framework ensures that data subjects' rights over their data are inalienable and independent of any terms and conditions. Since the right is available *in rem* (i.e., against the world at large), data controllers will be unable to use consent as a defence to escape liability. In fact, data controllers will be obliged to ensure that no rights of the data subject are violated as a consequence of the manner in which the data was collected or processed.

The data rights shall include the following:

a. *The Right to Fair Treatment*:

The data subject has the right to be treated without bias when a data controller arrives at a decision or makes a determination about the data subject by processing any data. As machine learning relies heavily on patterns to make decisions, it is easy for bias to seep in[8]. This results in patterns that produces consistent but sometimes discriminatory results.

For instance, a machine learning algorithm in a bank detects a pattern that, on the basis of the data that is already fed into it, and uses it to decide that loans must not be given to people from backward castes. A data subject from a backward caste who has good credit is unjustifiably denied a loan on the basis of this determination. By expressly providing for the right to fair treatment, the proposed statute will allow wronged data subjects to assert that they were harmed due to biased data processing.

b. *The Right to Information*:

The data subject has the right to information about all data pertaining to him that is possessed by or under the control of a data controller. The data subject must have the right to know what use any such data has been put to, the persons or entities with whom his data has been shared, and the persons or entities who have had access to his data. The data subject will also have the right to require that the data controller correct any errors or omissions in the data that are within the ability of the data controller to correct, provided such errors or omissions are verifiably true.

c. *The Right to a Data Security*:

The data subject shall have the right to be assured of the security of his data at all times. This means that the data controllers are obliged to store all data under their control securely. Further, the processing and transfer of such data must be in compliance with adequate security practices and procedures.

d. *The Right Against Processing*:

Data subjects who do not want their data to be processed shall have the right to require that the data controller stop processing their data forthwith. The statute could be implemented in such a manner as to offer the data subject granular control rather than presenting just binary options.

For instance, the statute could require the data controller to show the data subject all the ways in which his data is processed. Then, through an interactive dashboard, the data subject could granularly withdraw consent to the processing of certain categories of data. The interactive dashboard should be designed such that it informs the data subject of the consequences of withdrawing consent on the services provided (for instance, withdrawing consent to collect location data would mean that the data controller could no longer provide the services of turn by turn navigation).

In the proposed Rights Model, the data subject does not have to formally consent each time the data controller collects his data. Rather, the obligation is shifted to the data controller, who will need to ensure that data is processed or otherwise dealt with without violating any of the rights of the data subjects.

2. Data controllers will be liable for any provable harm to the data subject

The Rights Model will identify the types of harm that can result from a breach of any of the data rights. The data controller shall have the primary obligation to remediate the harm. This could include resetting the controller's record so that the data subject can be returned to the position he was in before the harm occurred. Alternatively, the data controller must be held liable for any loss suffered by the data subject due to improper data processing or a data breach.

In the event of a data breach, the controller will be obliged to issue a data breach notification. Where harm is the result of improper processing, the data controller must rectify the algorithm and notify data subjects who might have been affected.

If the data controller is to be held accountable for all consequences a data subject suffers, it is important to establish what constitutes harm under the new paradigm. While it is not possible to exhaustively enumerate all categories of harm caused by a privacy breach, the following types could be considered:

a. Financial harm:

The statute must recognise any direct or indirect financial harm that the data subject suffers as a consequence of processing of data, its transfer, or due to a security breach. While direct financial harm refers to financial losses, indirect financial harm is defined as any harm that can be quantified in financial terms even if it does not directly result in financial loss.

b. Reputational harm:

The statute must also recognise the intangible harm caused to the reputation or social standing of a data subject as a consequence of processing of data, its transfer, or due to a security breach. As a consequence of such harm, the data subject may find it hard to get a job, may be persecuted for criminal acts perpetrated in his name or otherwise shunned in society.

c. Harm due to Manipulation of Choice:

This type of harm is caused when the choices available to a data subject are limited on the basis of the data already provided. This limitation of choice results in the data subject having limited access to information and products or services.

The data right to fair treatment is infringed by this harm. To avoid this harm, the data processor must strike a balance between presenting the data subject with recommendations and completely denying the data subject access to any information or choices outside these recommendations.

d. *Harm due to Discrimination*:

This type of harm is caused when data is processed in an unfair and discriminatory manner against a data subject in terms of the products or services that such data subject is entitled to.

The data right to fair treatment is infringed by this harm as well. The data controller must ensure that processing data does not result in any discrimination. It should also ensure that the processing of data bears a demonstrable relationship to the purpose for which it was collected.

Whenever a harm occurs, the data controller's first responsibility is to remediate the harm caused and, if this is not possible, compensate the data subject for the harm caused. If a data breach has occurred (or is suspected to have occurred), then the data controller must issue a data breach notification within 24 hours of discovering the same. In case of harm due to improper processing, the data controller must rectify breaches and processing algorithms and notify the data subjects. The data controller should, additionally, take the necessary steps to reset the record so that the data subject can be returned to the position he was in before the harm occurred to her.

3. Liability of the data controllers should be severe, and tied to turnover

The data controller will be held liable for any harm caused to the data subject as a consequence of the breach of data rights. Where more than one data controller is involved in the processing that resulted in harm to the data subject, they will all be jointly and severally liable for the consequence. Data controllers should be subject to significant penalties to ensure that within this accountability framework, they take their fiduciary responsibility seriously.

In addition to compensating the data subject to the full extent of any loss caused due to breach, the data controller shall be obliged to take all steps necessary to reset the record so that the data subject can be returned to the position that existed prior to the occurrence of the harm.

Once the harm is identified, it should be possible to determine how it was caused. This would allow us to identify the data controller responsible for the harm and hold it liable to compensate the harmed data subject. The data controller should be accountable for any actions that, directly or indirectly, cause harm to the data subject. Indirect harm would include the transfer of data to a third party who subsequently deals with the data irresponsibly or otherwise violates the privacy of the data subject.

Since modern databases are inter-connected, it is possible that even without consciously transferring data, numerous data controllers have access to a given data set and therefore might have contributed (in varying degrees) to the harm caused to the data subject. In such events, separately apportioning the liability between these data controllers will be infeasible. In such scenarios, all the data controllers who may have been responsible shall be made jointly and severally liable for the consequences.

Data controllers must be aware that, under the accountability principle of the new framework, there are severe consequences for failing to meet their fiduciary obligations. The European Union's General Data Protection Regulations ("**GDPR**") stipulates that the penalties for failing to comply with data protection obligations is in the order of 5% of the global turnover of the data controller. In order to make an accountability based law truly effective in India, penalties of a similar magnitude must be included.

IMPLEMENTATION MECHANICS OF THE RIGHTS MODEL

1. Statutorily prescribed security standards

The Rights Model must spell out in sufficient detail the security obligations that data controllers need to comply with while handling data belonging to a data subject. These regulations must be responsive to changes in technology and the compliance by data controllers with these obligations should be appropriately auditable.

One of the principal obligations of the data controller under the Rights Model is to ensure that a data subject's right to data security is preserved. Data security should be addressed in terms of both *physical security* as well as *technical and operational security*. Sufficient work has been done around the world to articulate these standards and it would be appropriate to include references to those provisions within the new model.

The model should also have appropriate mechanisms to ensure that the security standards articulated are responsive to changing technology. To this end, the new model must include the creation of an advisory board that comprises persons skilled in their understanding of data security measures. These individuals will be tasked with ensuring that the standards are up to date. Compliance with such standards must be the subject of regular audit that should fall within the remit of the learned intermediaries (discussed subsequently).

2. Learned Intermediaries to audit and remedy data

A Learned Intermediary must be constituted under the Rights Model. The Learned Intermediary, performing the role of an auditor, will review the data processing algorithms of a data controller to evaluate whether they are privacy neutral.

This Document proposes 3 stages of audit– (i) Database Query Review; (ii)Black Box Audits; and (iii) Algorithm Review. The Learned Intermediaries should be obliged and capable of indicating appropriate remedial measures to neutralise any bias that they detect. The Rights

Model could also recommend the establishment of an open algorithm format that allows Learned Intermediaries to audit more effectively.

The fiduciary obligation that the Rights Model imposes on the data controller may not, of itself, be sufficient to ensure compliance. The extent of this fiduciary obligation can often be difficult to determine. Left solely in the hands of the data controller, it could be interpreted to suit the controller's interests.

Since harm to the data subject could take place over an extended period of time, there is a risk that it could remain undiscovered unless the actions of the data controller are actively monitored. Harms caused by machine learning algorithms are often the result of inherent algorithmic bias. This bias is often extremely difficult to detect and puts the very process of determining a breach of fiduciary obligations beyond the technical capabilities of most people.

In order to address this concern, an additional layer of supervision should be added to the framework. We propose that a class of learned intermediaries be created to ensure that the data subject's data rights are not violated. Learned Intermediaries will be technical personnel trained to evaluate the output of machine learning algorithms and detect bias on the margins. Learned Intermediaries will also be legitimate auditors of data controllers and the proposed law will lay down a mechanism for them to be certified. They must conduct periodic reviews of the data controller's algorithms with the objective of making them stronger and more privacy protective, and not just point out problems in its functioning. They should be capable of indicating appropriate remedial measures if they detect bias in an algorithm. For instance, a Learned Intermediary can introduce an appropriate amount of noise into the processing so that any bias caused over time due to a set pattern is fuzzed out.[9]

The Learned Intermediaries could adopt the following staged approach to review the operations of the data controller:

a. *Publication of Queries to Databases*:

The Rights Model must mandate all data controllers to publish the queries they make on databases containing information about the data subjects. As disclosure of this information will not detrimentally affect the data controllers' proprietary rights over their algorithms, it is easily realisable. In fact, enabling Learned Intermediaries to review the disclosed queries will help detect bad actors who are querying the database for information outside the purpose for which data has been collected.

For instance, in the context of loan processing, a query on a bank's database relating to the caste of loan applicants is clearly irrelevant to the purpose for which the data was processed.

b. *Black Box Audit*

In a black box audit, the actual algorithms of the data controllers are not reviewed. Instead, the audit compares the input algorithm to the resulting output to verify that the algorithm is in fact performing in a privacy preserving manner. This mechanism is designed to strike the balance between the auditability of the algorithm on one hand and the need to preserve proprietary advantage of the data controller on the other.

Data controllers should be mandated to make themselves and their algorithms accessible to the Learned Intermediaries for a black box audit.

c. *Access to Algorithms*

Should the data controller's algorithms appear to operate in a manner detrimental to the data subjects' rights, the data controller can be required to provide the Learned Intermediary with access to its algorithms. Where possible, open source algorithms could be used for providing such access, so that their workings are visible for inspection while at the same time the data controller's proprietary details of the algorithms are protected.

It is anticipated that in time, data subjects will flock to those data controllers whose algorithms are consistently certified by Learned Intermediaries as being privacy neutral. This will incentivise the data controller to align with the accountability principle in the Rights Model.

3. Regulation and primary adjudication through a Data Commissioner

The Rights Model should include the creation of a regulatory authority, such as a Data Commissioner who will be responsible for redressal of grievances as well as for establishment of technology responsive standards of accountability and transparency.

The Rights Model must establish a regulatory authority, such as a Data Commissioner, who will be responsible for the following:

a. *Establishing standards of accountability and transparency*

The Data Commissioner will have the broad responsibility of developing and implementing appropriate standards of accountability and transparency, and enforcing them among data controllers. These standards must be updated regularly (at least annually) to take into account advances in technology.

b. *Handling instances of data rights violations and providing redress*

In addition to the primary administrative responsibility outlined above, the Data Commissioner will also serve as the principal adjudicative functionary in the Rights Model.

Should a learned intermediary report that a data controller is processing data in a manner detrimental to a data subject, or if a data subject complains about harm caused to him as a result of improper processing, the Data Commissioner will be empowered to investigate the matter and make a determination as to whether or not such harm was caused. The Data Commissioner will have the power to pass appropriate orders to ensure that the harm is reversed or compensated for, as the case may be. In this regard, the damages awarded by the Data Commissioner can, under no circumstances, exceed 5% of the global turnover of the entity in question.

The Rights Model should establish technology enabled protocols for processing complaints and data subjects should be encouraged to use these online mechanisms to obtain redress. The use of digital means will allow a Data Commissioner to operate in a manner that allows greater flexibility and ensures the removal of bureaucracy. More thinking will be required in relation to the design of the regulator, the nature of its accountability, the governance framework within which the regulator will operate, and the process through which the regulator may make further regulations.

4. The Rights Model will be applicable to the State

The new privacy framework should apply to the government and other agencies of the State. The state should collect and process only that data that is relevant for its purpose. Even though there is no need for a national security and law enforcement exception in the absence of consent, the use of personal data for these purposes should be in line with the accountability framework. A special court could be constituted under this framework to authorise, where necessary, the use of broad data collection powers that would have otherwise violated the accountability obligations of the data controller.

It is proposed that the Rights Model apply equally to the government and agencies of the state as it does to private entities.

a. *Definition of Data Controller to encompass state agencies*

The term data controller should be defined to include any entity that collects information from a data subject, whether in the public sector or the private sector.

b. *Access to data for investigation, or in national interest*

Most data protection statutes around the world have exceptions built into them to allow law enforcement agencies to access personal information to aid criminal investigation, and also to allow security agencies to access this information in the interests of national security. These typically operate as exceptions to the principle of consent. Under the proposed model, there is no requirement to procure consent before accessing data about a data subject and to that extent, there is no need to expressly articulate this exception.

Instead, the Rights Model will specify that the government be obliged to ensure that no harm is caused to the concerned data subjects when it accesses their personal information (in the interests of national security or in the course of investigating a crime.) Express provisions should be added in the statute to ensure that investigations do not end up as fishing expeditions trampling over the personal privacy of a wide swathe of persons.

c. *Special court to conduct broad investigations on state actions affecting privacy*

As mentioned above, the State may seek exceptions to breach the data rights of data subjects in certain extenuating circumstances (for example, in the interest of national security). In order to ensure that the government protects the data rights of its residents even under these circumstances, it might be useful to establish a special court (along the lines of the court under the Foreign Intelligence Surveillance Act in the USA) (the **Special Court**).

Broad investigations affecting personal privacy must be conducted by the State only after obtaining prior approval from the Special Court. The Special Court must assess whether the pursuit of a line of investigation that could potentially affect the data rights of a wide range of data subjects is justified or not. In arriving at that conclusion, the Special Court will evaluate whether the data sought to be collected suits the purpose and whether collection is limited to what is necessary to fulfil that purpose. It could also satisfy itself about which agencies and third parties the data will be transferred to. The Data Commissioner could be an *ex officio* member of this court.

5. Consent to continue having importance

*Data controllers can ask data subjects to consent to the terms of their privacy policies, but cannot rely on such consent as an excuse for failing to meet accountability obligations***.**

Nothing contained here should be misconstrued to mean that under this new construct, data controllers are prohibited from seeking prior consent from data subjects. In most cases, data controllers who operate across multiple jurisdictions are obliged, under the laws of those countries, to obtain consent prior to data collection and processing. If India were to adopt an accountability based privacy framework, it would be one of the only countries in the world that dispenses with consent entirely. It is therefore proposed that data controllers be free to continue to collect consent using the same model they use around the world on the clear understanding that any consent so collected will not operate as an indemnity that will diminish their obligation to remain accountable under the new data protection framework.

In other words, while data controllers will not be prohibited from continuing to ask data subjects to consent to the terms of their privacy policies, they will not be able to rely on such consent as a reason for any violation of the data rights of a data subject or an excuse for not meeting their accountability obligations.

CONSIDERATIONS

We note that the proposed Rights Model gives rise to some questions about the scope of the new statute, the extent of rights and responsibilities under it, etc. In this section, we address some of these concerns:

1. **On why we need a new model of data protection**

   a. Under the current legal framework, it is impossible to determine the full extent to which a person's rights over their data are inalienable. Inalienability should be born out of a statute. For this reason, the data right to privacy (and other connected data rights) must be codified in a law.

   a. Further, the lack of such a legislation will make it easier to mould a law on a rights-based approach rather than a consent based approach.

   b. The law will deal with the data rights that individuals have and the responsibility that data controllers have with regard to these rights. It will not concern itself with the manner in which the data collected can be monetised.

2. **On whether the Rights Model will clash with the Constitutional Right to Privacy?**

   a. While it has not yet been established that the right to privacy is an inalienable constitutional right, the Supreme Court is expected to rule on this issue in the foreseeable future. Nonetheless, it is anticipated that the Supreme Court's decision will not materially impact the model proposed here.

   b. While the Supreme Court is concerned with the right to privacy in general, a data protection statute (as addressed under the Rights Model) will be more limited to privacy as it pertains to personal data. However, this allows the statute on data protection to be more detailed than one on the right to privacy, which is more generic.

   c. On the other hand, the right to privacy and the statutory data rights proposed here share some common underlying principles (for instance, autonomy and security). To this extent, the Supreme Court's ruling may have some repercussions on the proposed Rights Model.

3. **On Data Rights and the extent to which the Rights Model will help define them**

   a. The rights enumerated in the statute must be as simple as possible to avoid complication in interpretation in the future.

   b. The autonomy principle allows a data subject to request that his data to no longer be used by a data controller. However, requesting the data controller to stop using a person's data will mean that the person can no longer avail the services / product offered by the controller. While this may work in cases where the data controller is a private entity, it is unclear what happens where the state collects / processes individual data for authentication purposes.

4. **On the responsibilities of data controllers with deeper pockets**

   a. Attaching joint and several liability on data controllers will generate benefits for data subjects. It will ensure that the data controller with the deepest pockets will shoulder the highest burden from amongst the data controllers.

   b. Civil liability under the Rights Model is tied to the data controller's annual global turnover. The prospect of high penalty is intended to act as a deterrence. The deeper the data controller's pockets, the higher its potential punishment and so, the higher is the likelihood that it will be responsible in its data practices.

5. **On data controllers being allowed to limit their liability**

   a. Data controllers may seek to limit their liability under the terms and conditions with the data subjects. They could frame these terms and conditions to include the processing of data by technological means such as machine learning.

   b. However, the Rights Model stresses that the rights of the data subjects are more susceptible to harm. Therefore, must ensure that data controllers are dissuaded from limiting their liability to unrealistically low sums.

   c. Currently, the terms and conditions / privacy agreement that each data controller makes his data subjects sign overrides legal provisions. The statute born out of the Rights Model will make any agreement between controllers and subjects to be subject to its provisions. This statute, on its part, should clarify that there is no limited liability through the obtaining of consent.

6. **On Learned Intermediaries: Scope of Responsibility, Certification, Extent of Discretion**

   a. Learned Intermediaries may not be needed to audit all data controllers. The statute must specify that audits by a Learned Intermediary will be mandatory for a company (or other kind of data controller) above a defined threshold, while, for start-ups and smaller entities, Learned Intermediary audits may be recommended but not compulsory.

   b. Equally, if the business of a company / data processor is such that the consequences of harm are more real / huge, then audits by Learned Intermediaries must be mandatory.

   c. While the Learned Intermediary is tasked with the responsibility to detecting and undoing, there is no standard laying down what constitutes bias and what does not. Therefore, the Learned Intermediary will have discretion to decide what constitutes bias, and this discretion will be checked by the Data Commissioner from time to time.

d. Also, the new statute must ensure that Learned Intermediaries are certified by the appropriate authority to perform their functions. This will aid in reducing the margin of error.

e. There may also be merit to the Learned Intermediary's findings being published as ratings, along the lines of credit ratings published by credit rating agencies.

FURTHER WORK

Further ideation / structuring is needed on the following issues:

1. Biases: Incremental nature and pre-emption of harm

a. Algorithms have biases built in. Machine learning enables computers to detect patterns and make automated decisions on the basis of these patterns. Sometimes, the patterns formed out of such machine learning is based on biased data (for example, a data input that shows black single mothers being unable to pay loans becomes part of a pattern that subsequently denies giving loans to otherwise creditworthy black single mothers). This raises the question: Is a data point categorized as biased right at input? If yes, then what data are tagged biased, and what are allowed?

b. Biases are also incremental. This means that the patterning can stay latent for a long time and be detected years later. In the interim, large numbers of people can be affected. How would this be remedied?

c. Additionally, the functions that Learned Intermediaries perform are corrective in nature. Their ability to conduct black box audits and algorithm reviews are post-default measures. How effectively will these Intermediaries be able to prevent biases in algorithms?

2. Monetisation of data by the State

a. The State, which collects large chunks of information, stands to gain financially from transferring aggregates of personal data it has collected to third party entities. How will this be dealt with?

b. A possible counter-argument at this point, is that the data subject's data will likely never be accessed at an individual level, but is aggregated with similar data. How long this will last is uncertain.

[1] Rahul Matthan is a lawyer specialising in Telecommunications, Media and Technology law and is a partner at Trilegal. He is also a Fellow of the Technology and Policy Programme at the Takshashila Institution. This document is edited by Manasa Venkataraman and Ajay Patri, the Takshashila Institution. This document is prepared for the purposes of discussion and debate only and does not necessarily constitute Takshashila's policy recommendations. To contact us about the research write to research@takshashila.org.in or visit takshashila.org.in.

[2] Regulating the Internet of Things: First Steps Toward Managing Discrimination, Privacy, Security & Consent, Texas Law Review, 1 March 2014 (last visited 21 April 2017); Big data: Credit where Credit's Due, Financial Times, 5 February 2015 (last visited 21 April 2017); Adam D. Thierer, *The Internet of Things and Wearable Technology: Addressing Privacy and Security Concerns without Derailing Innovation*, Vol. 21 Richmond Journal of Law and Technology, p. 6, 18 February 2015.

[3] Jonathan A. Obar and Anne Oeldorf-Hirsch, The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services, TPRC 44: The 44th Research Conference on Communication, Information and Internet Policy 2016, 24 August 2016 (last visited 21 April 2017).

[4] Aaron Smith, Half of Online Americans Don't Know what a Privacy Policy Is, Pew Research Center, 4 December 2014 (last visited 21 April 2017).

[5] Aleecia M. McDonald and Lorrie Faith Cranor, The Cost of Reading Privacy Policies, I/S: A Journal of Law and Policy for the Information Society,2008 (last visited 21 April 2017).

[6] Learning Hannah Devlin, Discrimination by Algorithm: Scientists Devise Test to Detect AI Bias, The Guardian, 19 December 2016 (last visited 21 April 2017).

[7] Automated high frequency algorithmic trading has been the cause of at least three stock market crashes in the United States of America since 2010. JP Buntinx, Top 3 Financial Crashes Caused by High Frequency Trading Algorithms, The Merkle, 28 February 2017 available at < https://themerkle.com/top-3-financial-crashes-caused-by-high-frequency-trading-algorithms/> (last visited on 19 July 2017)

[8] Aylin Caliskan, Joanna J. Bryson and Arvind Narayanan, Semantics derived automatically from language corpora contain human-like biases, Science Magazine, 9 July 2017, (last visited 9 July 2017).

[9] Stephanie Pappas, Bad news: artificial intelligence is racist, too, Live Science, 13 April 2017, (last visited 9 July 2017); Jennifer Smith, UW-Madison scholars tackle bias in algorithms, UW-Madison News, 3 July 2017, (last visited 9 July 2017).