



2200 PENNSYLVANIA AVENUE NW
WASHINGTON, DC 20037
TEL 202 • 955 • 1500
FAX 202 • 778 • 2201

VIA ELECTRONIC FILING

October 8, 2014

Federal Trade Commission, Office of the Secretary
Room H-113 (Annex X)
600 Pennsylvania Avenue NW
Washington, DC 20580

Re: Comment of the Centre for Information Policy Leadership on “Big Data: A Tool for Inclusion or Exclusion? Workshop, Project No. P145406”

The Centre for Information Policy Leadership (the “Centre”)¹ appreciates the opportunity to respond to the Federal Trade Commission’s request for comments on issues related to the FTC’s recent workshop on “Big Data: A Tool for Inclusion or Exclusion?” The Centre strongly supports the attention given by the FTC to the opportunities and challenges of big data analytics. As an organization whose members are global businesses committed to using personal data responsibly, the Centre, too, is significantly focused on developing solutions to the challenges of big data, so that businesses and society can safely reap its benefits.

I. Introduction

The essential message of the FTC’s September 15, 2014 workshop “Big Data: A Tool for Inclusion or Exclusion?” was that big data analytics bring enormous benefits to society, but they also have the potential to wrongfully label individuals, stigmatize protected classes, and perpetuate societal prejudices and discrimination. On the benefits side, the many examples included predicting flu outbreaks based on online searches or targeting advertisements only to those who are likely to want them, and uncovering discrimination through big data analysis. Examples of harmful effects of big data analytics ranged from a plethora of speculative scenarios, such as the ability to draw negative inferences about the economic status of a

¹ The Centre for Information Policy Leadership is a privacy and data protection think tank in the law firm of Hunton & Williams LLP and is financially supported by approximately 40 member companies that are leaders in key sectors of the global economy. The Centre’s mission is to engage in thought leadership and develop best practices to ensure effective privacy protection in the modern information age. For more information, please see the Centre’s website at <http://www.informationpolicycentre.com/>. Note that nothing in this submission should be construed as representing the views of any individual Centre member or of the law firm Hunton & Williams.

consumer from the fact that she buys vegetable seeds, to making decisions based on “proxy” variables that are statistically associated with protected minorities, thereby inadvertently or surreptitiously discriminating against these minorities, to the well-known, real-life example of the “pot-hole finding” mobile app results that were skewed in favor of wealthier Boston neighborhoods — those with greater smartphone penetration — until the data was complemented by gathering additional data through other means.

One speaker called the range of opportunity and detriment of big data the “data paradox.” Data can be used in ways that help or hurt. While the speakers discussed a range of possible responses to deal with this paradox, including taking a “wait and see” approach and relying on existing credit, finance, employment and consumer protection laws that already cover many big data risks, several speakers suggested that the sheer variety of big data applications precluded a “one size fits all” approach, or they pointed out that protecting from harm in this context is “situational” and that “framework models” for addressing such context-specific risks and harms should be devised.

The Centre agrees with the speakers who called for a situational and contextual approach to the risks and harms of big data. And, as several speakers noted, one such approach — the risk-based approach to privacy (also referred to as “benefits-risk” analysis) — already exists. Risk assessments are already being used by organizations as part of their “organizational accountability” and privacy compliance programs. In some important cases, they are already part of the law. See the “legitimate interest” ground for data processing in the EU Data Protection Directive², mandatory Privacy Impact Assessments in an increasing number of jurisdictions, and the FTC’s unfairness authority³, for example.

Organizational accountability programs that incorporate an effective risk-based approach and appropriate transparency are uniquely suitable for addressing context-specific challenges of big data. As such, the Centre urges the FTC to further focus on and incorporate these concepts as it develops its approach to big data.

II. Accountability, the Risk-based Approach to Privacy, and Transparency

The recent dramatic changes in the information economy in terms of complexity, volume, speed and analytics-based insight (encapsulated in the term “big data”) demand adjustments to the traditional paradigms for protecting consumers from privacy harms and other harms related to information use, as well as further refinement of existing mechanisms and development of new

² Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data, available at <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:en:HTML>; see also Article 29 Data Protection Working Party, Opinion 06/2014 on the notion of legitimate interest of the data controller under Article 7 of Directive 95/46/EC, adopted on 9 April 2014.

³ Section 5 of the FTC Act, 15 USC 45, prohibits “unfair and deceptive acts or practices in or affecting commerce.” Under FTC jurisprudence, an act or practice is unfair where it (i) causes or is likely to cause substantial consumer injury; (ii) cannot be reasonably avoided by consumers; and (iii) is not outweighed by countervailing benefits to consumers or to competition. Thus, an unfairness claim requires an assessment and weighing of the risks and benefits of the conduct at issue similar to the risk-based approach to privacy.

tools that can deliver effective protections even as technology and business models continue to change. These mechanisms and tools are organizational accountability, a risk-based approach to privacy and transparency.

A. Accountability

The value of creating organizational accountability through comprehensive corporate privacy compliance programs is now widely accepted by responsible businesses, not least due to the FTC's efforts through its enforcement actions and consent orders. Such comprehensive programs are essential for delivering effective privacy protections in the context of our information-based economy. They facilitate responsible and accountable information practices that protect individuals from harmful impacts caused by the use of personal information. They enable legal compliance, provide the framework for demonstrating such compliance (both internally to corporate boards and management, and to regulators), reduce commercial and reputational risks, and preserve shareholder value. They also help shift the burden of privacy protection from consumers to the organizations.⁴

Such organizational accountability programs also include enforceable industry codes of conduct and similar self-regulatory or co-regulatory mechanisms as envisioned, for example, in the 2012 Privacy Blueprint of the Administration⁵, and as already implemented in, for example, the APEC Cross-Border Privacy Rules (CBPR) system. Indeed, in addition to furthering organizational accountability by requiring comprehensive privacy compliance programs in its consent orders, the FTC has been supportive of enforceable self-regulation by industry.⁶

To make such corporate accountability schemes even more effective in the context of big data analytics, they should be augmented with effective risk assessments and improved transparency, as outlined below. Together, these measures would help accountable organizations interpret and implement existing privacy principles in the context of big data more effectively and deliver robust protections for individuals where collection and use of information cannot always be based on individual control and consent.

⁴ The Centre's previous pioneering work on accountability discusses in greater details the elements of accountability and corporate privacy programs. See <http://www.informationpolicycentre.com/accountability-based-privacy-governance/> for information and relevant white papers.

⁵ *Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Economy* (the "Privacy Blueprint"), available at <http://www.whitehouse.gov/sites/default/files/privacy-final.pdf>.

⁶ See, e.g., *Protecting Consumer Privacy in an Era of Rapid Change, Recommendations for Businesses and Policymakers*, FTC Report (Privacy Report), March 2012, available at: <http://www.ftc.gov/reports/protecting-consumer-privacy-era-rapid-change-recommendations-businesses-policymakers>.

B. Risk Assessments and the Risk-based Approach

The risk-based approach is closely linked with the principle of accountability. This is because risk assessment is one of the essential elements of accountability and corporate privacy programs.

However, risk-assessments traditionally focused primarily on the risks to the organization rather than individuals and society. Building on its earlier work on organizational accountability, the Centre recently launched a multiyear project to develop a comprehensive analytical framework for risk assessments that focus on the risks to individuals and society — the Privacy Risk Framework Project. Specifically, the project seeks to build consensus in collaboration with international privacy regulators, privacy experts and industry members, on what is meant by such privacy risks, and to create a practical framework and the tools for organizations to identify, quantify, prioritize and mitigate such risks. It also seeks to explore how regulators, such as the FTC, can employ a risk-based approach to enforcement that is consistent with the private sector's use of risk assessments within the context of organizational privacy and compliance programs. In June 2014, the Centre published its first white paper on the subject “A Risk-based Approach to Privacy: Improving Effectiveness in Practice.” (Attached in the Appendix).

Risk assessments can help accountable organizations determine whether and how to proceed with proposed information uses based on the potential risks and harms they may cause to individuals and, possibly, society. Risk assessments are an essential part of devising proper information security measures and integral to privacy by design. While they should be performed for all data uses, they are uniquely suited to enable responsible decisions in the context of big data for the following reasons:

- a. Understanding the likelihood and potential severity of harms to individuals that may result from proposed information uses in specific contexts allows organizations to devise appropriate and targeted mitigations and controls. It also facilitates weighing any residual risk of harms after mitigations have been implemented against the countervailing benefits of the proposed use, before making any decision with respect to such use. Such analysis would be relevant, for example, where an organization sought to avoid an FTC unfairness action.
- b. Risk assessments place the burden of privacy protection on the organization. They are especially useful in situations where individual control and consent would be too burdensome due to the absence of direct interaction with the individual (*e.g.*, if the data has been de-identified) or the complexity of the involved information processing, or if consent is otherwise impossible.⁷

⁷ Use of risk assessments in lieu of consent is not new. In fact, even the more restrictive European data protection law includes a similar concept that allows for the processing of data where consent is not feasible: the “legitimate interest ground” for data processing. *See, infra*, at n. 2.

Under that provision, organizations may collect, use or share data when it is in their “legitimate interest” to do so, provided such collection, use or sharing does not prejudice individuals’ rights and freedoms. The test for determining whether an organization may proceed on the basis of such legitimate interest requires considering the

- c. Risk assessments reduce inefficient deployment of organizational resources by allowing organizations to prioritize their controls according to the likelihood and severity of harm associated with a proposed data use. Such prioritization is likely to contribute to the overall effectiveness of privacy protections or other measures designed to protect consumers.
- d. Risk assessments based on a coherent methodology provide a useful tool to help interpret and apply existing privacy principles and legal requirements in the big data world. They are critical tools for calibrating compliance to the levels of risk and harms in a particular context and to enable the privacy principles to remain flexible and stand the test of time.

A crucial issue is how to identify and agree on the nature, classification, and quantification of risks. To yield effective protection, companies' risk-based approach must take an inclusive approach to harm. It must not only seek to identify and evaluate tangible harms such as bodily injury, financial and other economic harms, but must also consider intangible harms (whether cognizable under law or not), such as reputational harm, embarrassment, discrimination and stigmatization, as further described in the Appendix. As part of its ongoing Privacy Risk Framework Project, the Centre is currently working to develop consensus on the types of the harms that should be considered and how to classify, quantify and address them.

Furthermore, to be effective, risk assessments must consider the entire life cycle of information from collection, use and sharing, to disposal. They must also occur repeatedly over time, to evaluate and re-evaluate data uses on an ongoing basis. As such, the risk-based approach allows for a holistic approach to privacy protection that goes beyond the point of data collection and ensures that individuals continue to enjoy protections when their data are subsequently used and shared for various purposes in the complex data ecosystems of the modern information age.

Because of its significant potential in addressing the challenges of big data, privacy risk analysis and the risk-based approach should be incorporated in any future solutions to challenges that might be associated with big data analytics.⁸

impacts of the proposed data processing on the individual, and a balancing of the respective rights and interests of the organization and the individual. Thus, the European legitimate interest ground for data processing essentially envisages a risk-analysis consistent with the risk-based approach described herein. It is a noteworthy illustration of the importance of providing for a risk-based approach to privacy protection in appropriate circumstances. Its benefits include facilitating data collection, use, sharing and disclosure in circumstances where consent is not feasible, practicable or effective; enabling new uses of information for new purposes (such as big data and analytics), beyond the original purposes at the time of collection, provided such uses are not harmful to consumers and appropriate safeguards are implemented; and ensuring the protection of individuals' privacy, while allowing organizations to pursue the benefits of new technologies, products and services.

⁸ De-identification of personal information is also an essential tool for protecting and maximizing the benefits of big data. It is one of the mitigations and controls that can follow a risk analysis in the context of big data analytics and, as such, is an important element of the risk-based approach. To be as effective as possible, de-identification should combine sound technical de-identification processes with an organizational commitment and/or legal or contractual requirements not to re-identify the data, as proposed in the FTC's 2012 Privacy Report.

C. Transparency

Transparency is a cornerstone of the FTC's Privacy Framework as outlined in its 2012 Privacy Report. The Centre submits that it will become even more important in the context of big data where the sheer complexity and multiplicity of data uses will continue to eclipse the relevance of traditional legalistic or just-in-time notices, but where information use practices will, nevertheless, need to be communicated to consumers.

In situations where just-in-time notice is not feasible and information use decisions, therefore, cannot be based on obtaining individual consent, the concept of "notice" should be de-coupled from consent and be evolved into a broader concept of "transparency." This transparency would provide consumers reasonable access to appropriate general information about how their personal information may be used, the benefits of such use, the known potential impacts, and the measures undertaken to protect against harm, along with any other relevant information. This may be especially important where such uses go beyond the original context of the collection and where future uses are unusual, complex or not expected. A broadened concept of transparency, coupled with actual protections stemming from robust risk assessments within the context of accountable corporate privacy programs, can significantly contribute to creating consumer confidence and trust in contexts where just-in-time notice and individual consent are not feasible.

Some accountable organizations are already exploring ways to develop enhanced and modernized transparency models suitable for the big data context, (for example through privacy dashboards), that can deliver relevant information and even provide appropriate information management tools to consumers. As the FTC continues to develop its approach to big data, it should continue to emphasize transparency to consumers as an important component as well.

III. Conclusion

As the nation's principal privacy and consumer protection authority, the FTC can help ensure that big data can grow safely, including where traditional notice and consent are impractical, impossible or illusory. Well-designed, demonstrable and verifiable accountability programs, a risk-based approach to privacy, and transparency, are tools that can help deliver appropriate protections in such circumstances.

Additional work, however, to refine these tools and to socialize the risk methodology for the big data context needs to be done. That work is underway in many corners. Indeed, the FTC's own approach to privacy protection and its enforcement tools under the FTC Act, most notably its "unfairness authority," either are consistent with a risk-based approach or already rely on the approach directly. As such, the Centre encourages the FTC to further explore the potential of organizational accountability, the risk-based approach to privacy, and transparency, specifically to address the challenges of big data.

Thank you for accepting and considering our comments and recommendations. If you have any questions about this submission, please contact Bojana Bellamy, President, Centre for

Information Policy Leadership, bbellamy@hunton.com, or Markus Heyder, Vice President and Senior Policy Counselor, Centre for Information Policy Leadership, mheyder@hunton.com.

Appendix
A Risk-based Approach to Privacy:
Improving Effectiveness in Practice

A Risk-based Approach to Privacy: Improving Effectiveness in Practice

In January 2014, the Centre for Information Policy Leadership (the Centre) launched a multi-year project on the risk-based approach to privacy: The Privacy Risk Framework Project. This project elaborates on the Centre's earlier project on organisational accountability, particularly in seeking to develop the analytical framework and tools needed to implement certain key aspects of accountability. Specifically, the goals of this project are set forth in the following Project Vision Statement:

Principle-based data privacy laws often leave room for interpretation, leaving it both to organisations to make appropriate decisions on how to implement these principles and to regulators on how to interpret and enforce the law. The Privacy Risk Framework Project aims to bridge the gap between high-level privacy principles on one hand, and compliance on the ground on the other, by developing a methodology for organisations to apply, calibrate and implement abstract privacy obligations based on the actual risks and benefits of the proposed data processing. While certain types of risk assessments are already an integral part of accountable organisations' privacy management programs, they require further development. This project seeks to build consensus on what is meant by privacy risks to individuals (and society) and to create a practical framework to identify, prioritise and mitigate such risks so that principle-based privacy obligations can be implemented appropriately and effectively.

On March 20, 2014, the Centre held a workshop in Paris during which more than 50 privacy experts, industry representatives and regulators discussed their experiences and views with respect to the risk-based approach to privacy, the privacy risk framework and methodology, as well as goals and next steps in this project. This paper, titled "*A Risk-based Approach to Privacy: Improving Effectiveness in Practice*", is a developed version of the earlier discussion paper distributed to the participants of the workshop. It incorporates feedback from the Paris workshop and input received in subsequent consultations with Centre members and project participants.

I. Scope and Objectives

1. Data protection and privacy laws are meant to protect people, not data. But from what exactly are people being protected? What threats? What harms? What risks?

“Harm” in this paper is not meant to be a technical term. It simply signifies any damage, injury or negative impact—whether tangible or intangible, economic, non-economic or reputational—to an individual that may flow from the processing of personal data. It extends to any denial of fundamental rights and freedoms. The Privacy Risk Framework Project will explore whether and (if so) how it should also extend to any harms to society at large.

2. At a time when the new information age challenges accepted privacy concepts and practices and strains our limited compliance and enforcement resources, organisations and regulators need to prioritise their activities and find new ways to turn abstract requirements into real and effective privacy protections. In Europe, the notion of data protection as a fundamental human right has been reaffirmed by the European Court of Justice. In other parts of the world, high-level privacy principles continue to be articulated by policy-makers, legislators, courts and commentators. Yet, it is no longer enough—or sufficiently meaningful—to say solely that privacy is a human right and that the laws exist to safeguard “fundamental rights and freedoms”, nor that they are confined solely to existing privacy principles or fair information practices. New times call for new clarity and new pragmatism. A “Risk-based Approach to Privacy” can help deliver greater clarity and more effective data protection on the ground.
3. The risk-based approach goes beyond mere compliance with regulatory requirements. It goes to the heart of what responsible and accountable organisations seek to achieve, how they implement privacy requirements on the ground and how they demonstrate compliance. The risk-based approach may also help to clarify and communicate the underlying rationales for regulation.
4. As the pace of technological change outstrips the conventional thinking of law-makers, regulators and businesses, it is suggested that a calibrated, risk-based approach may improve the ability of businesses to take a better-informed and better-structured approach to the handling of colossal volumes of personal information that they collect, receive, store, use and share on a daily basis. These issues become more pressing as a greater number of companies seek to design, implement and demonstrate accountability through corporate privacy management programs and an ethical approach, often through programs of corporate and social responsibility. Increasingly, businesses, and their executives and boards need reassurances that their corporate programs are effective, and that they deliver required outcomes, both for the organisations themselves and for the individuals they seek to protect.
5. If the data privacy implications of products, services and other activities can be assessed from the perspective of their impact on individuals, can the likelihood of serious harm be reduced? Can the results of such assessments be reflected in better-targeted privacy programs and other safeguards? Also, how can it be made easier for non-experts to understand what they should—and should not—be

doing? How can privacy officers effectively communicate “the do’s and don’ts” of data privacy to an increasingly disparate audience of technologists, data scientists, privacy engineers and business leaders within their organisations?

6. Could a new consensus on a risk-based approach also help regulators fix and communicate their priorities for interpreting and enforcing the rules? Could it also give businesses more predictability and a better idea of what to expect and how best to avoid regulatory trouble?
7. In the longer term, how might this approach help policy-makers and legislators shape rules for the future that are more effective, less burdensome on businesses and individuals and take into account more precisely the risks to individuals and to the well-being of society, but without disenfranchising the individual?
8. The Centre’s Risk Project follows up on our pioneering work on accountability over the past five years. The project seeks to answer some of these questions and to explore the benefits of taking a more “Risk-based Approach to Privacy”. Specifically, this initial paper sets out issues and key learnings so far, with a first attempt to develop a framework to improve the ability of businesses to understand, identify, assess and manage privacy risks. This framework would also improve organisations’ ability to demonstrate to a third party, including a regulator, their “accountability” by enabling them to show specifically how and why they have reached certain data processing decisions.

II. Emerging Thinking

9. A number of headline messages have started to emerge from various workshops and discussions held in the last couple of years on, or around, the scope for a more risk-based approach.¹
10. In summary, the key messages and findings so far are as follows:
 - A risk-based approach is worth exploring for several reasons, all ultimately focused on improving the effectiveness of privacy protections in practice.
 - A risk-based approach should largely build on existing and emerging legislative provisions which already require consideration of privacy risks to individuals.

¹ In addition to the Centre’s previously mentioned Risk workshop in Paris (*see* p. 1), these included the Centre’s Accountability project workshops in Warsaw and Toronto, a session on risk at the 35th International Conference of Data Protection and Privacy Commissioners in Warsaw and an informal workshop sponsored by The Privacy Projects in London. The reports of the “Data Use and Impact Global Workshop” and the “Data Protection Principles for the 21st Century” have both drawn attention to the need for greater focus on the risks attendant on the various uses of data. A risk discussion also features heavily in the May 2014 white paper of the World Economic Forum entitled “Rethinking Personal Data: A New Lens for Strengthening Trust”. Further, on May 30, 2014, the Article 29 Data Protection Working Party adopted a “Statement on the role of a risk-based approach in data protection legal frameworks”.

- The risk-based approach is not meant to replace or negate existing privacy regulation and data protection principles. The approach and risk framework methodology primarily aim to:
 - a) complement the existing laws and regulations and facilitate the application of existing data protection principles and requirements;
 - b) help implement the existing legal requirements and privacy principles in a particular context, with greater flexibility and more agility that is required in the new information age, by taking into account the risks to individuals; and
 - c) improve the delivery of effective data protection in practice—benefitting individuals and organisations seeking more effective, systematic and demonstrable compliance.
- This means, in particular, providing clearer steers for accountable and responsible organisations that seek to “get it right” by preventing problems, often by going beyond compliance with legal requirements and regulators’ expectations. This may be for reputational, commercial or other reasons of enlightened self-interest.
- A risk-based approach has considerable potential to interpret, elaborate and make meaningful requirements and fundamental data protection principles which inevitably are often cast in general terms. Here, it is especially important to meet the growing needs of non-experts in privacy or data protection —engineers, data scientists, clinicians and many others—who need to grapple with these requirements and principles and reflect on the prospective impact of the new technologies and services they are developing.
- While the risk-based approach may be used to calibrate obligations and compliance of organisations, it should not be seen as a dilution of individuals’ rights, nor as a means of avoiding legal obligations.
- A risk-based approach is closely linked to the setting of priorities: “Selective to be Effective”. It helps organisations and regulators to concentrate on what really matters and to avoid wasting scarce resources on less important or bureaucratic requirements that neither benefit individuals nor better protect their information.
- The primary focus should be on **significant privacy risks** for individuals. In other words, in a given situation, the question should be whether there is **a significant likelihood that an identified threat could lead to a recognised harm with a significant degree of seriousness.**
- There is a particular benefit in developing a common and objective approach to risk management and an objective notion of harm or adverse impact to individuals that are acceptable and useful to as many businesses and regulators as possible.

- A similar approach might be applied to assessing risks and harms to society, although whether organisations can or should assess societal harms may require further consideration.
- Attempts to manage privacy risks should be integrated as closely as possible alongside well-established risk management processes
- but any approach must be kept as simple as possible and should be meaningful to SMEs (small and medium enterprises) and individuals as well as to large businesses, public bodies and regulators.
- As a risk-based approach will usually take the organisation beyond legal compliance in particular jurisdictions, it could be used as a tool to build and implement a consistent global program focused on the real priorities. More ambitiously, there is scope to improve the prospects for global inter-operability because following a common and consistent methodology to risk assessment would create harmonised practices and outcomes and, in turn, improve trust among regulators and individuals in different jurisdictions. It would also improve the ability of privacy authorities to cooperate on enforcement across borders.
- Any attempt to assess and manage risks in terms of impact on individuals and society would be novel. Hitherto, very few organisations or regulators have taken this as their rationale or motivation. Any structured encouragement for organisations to think in advance about the potentially negative impact of new developments should be welcome.
- Unsurprisingly, there is little agreement on what is meant by the “privacy risks” faced by individuals and society. The identification and classification of privacy risks must be settled before continuing work on how best to address them in a structured way.
- As a starting point, initial consensus on the nature of “privacy risks”, in terms of the **threats** and **harms**, would be useful, together with agreed methodologies for assessing **likelihood** and **seriousness** and balancing the results against the **benefits**.

III. Threats

11. When assessing threats, it is important to consider a whole lifecycle of information and data processing. Some threats will be visible at the time of collection, but some will emerge later, during the use or disclosure of data. It is important to note that the threats may also change during the lifecycle of information—old threats may disappear and new ones may become prominent.
12. Threats usually arise from processing personal data, which does or could relate to an identifiable individual. As anonymisation, however, becomes less absolute, all forms of data should be seen as capable of presenting privacy risks.

13. A wide approach to the threats arising over the lifecycle of data should therefore include both *activities* and *characteristics*. It is suggested that the following should be considered as the threats arising from data processing:
- unjustifiable or excessive collection of data;
 - use or storage of inaccurate or outdated data;
 - inappropriate use of data, including:
 - a) use of data beyond individuals' reasonable expectations;
 - b) unusual use of data beyond societal norms, where any reasonable individual in this context would object; or
 - c) unjustifiable inference or decision-making, which the organisation cannot objectively defend;
 - lost or stolen data; and
 - unjustifiable or unauthorised access, transfer, sharing or publishing of data.
14. In each case of the above threats, objective judgments will be needed about the a) likelihood of a threat causing harm to individuals, and b) the severity of that impact if it materialises. This means that the assessment of a threat arising from data processing must always be *contextual*. In other words, a flexibility is required that recognises context as an important factor in determining the level of threat and its potential to cause harm. In a risk-based environment, it is the *use* (including disclosure) of the information that arguably poses the greatest threat and where particular attention must be focused. This also has the advantage of avoiding the familiar practical problems of over-emphasis on collection solely, and of over-reliance on legalistic notice and consent, that result in information overload for individuals. Finally, this approach is also helpful in situations where there is no interface with individuals, or where the data are not collected directly from them.
15. Accordingly, neither information notification to the individual nor consent are by themselves a panacea. A use of personal information may be inappropriate or create significant privacy risks even though that use may have been specified or foreseen. The prominence and the extent of the individual's freedom of choice will be amongst the factors to consider, and may play a part in conditioning expectations, but neither "small print" disclosure, nor apparent consent, can, by themselves, justify an "unusual" use.

IV. Harms

16. There are three types of harm² that any of the identified threats could present:
- tangible damage to individuals;

² See explanation of the term "harm" on page 2.

- intangible distress to individuals; and
- societal.

17. **Tangible damage**, normally physical or economic, includes:

- bodily harm;
- loss of liberty or freedom of movement;
- damage to earning power; and
- other significant damage to economic interests, for example arising from identity theft.

18. **Intangible distress**, assessed objectively, includes:

- detriment arising from monitoring or exposure of identity, characteristics, activity, associations or opinions;
- chilling effect on freedom of speech, association, etc.;
- reputational harm;
- personal, family, workplace or social fear, embarrassment, apprehension or anxiety;
- unacceptable intrusion into private life; and
- discrimination or stigmatisation.

19. For both tangible damage and intangible distress, the harm may be potential (it could or would have this effect) or actual (it will, is having or has had this effect).

20. While risk assessment involves tests of foreseeability, these must be objective descriptors of harm—it is harm imposed on the reasonable man or woman in this context. In the same way as tort law ignores the “egg-shell skull”, the test is not, and cannot be, concerned with the impact on each particular individual, let alone an individual with particular sensibilities. Finally, the test must again be context-driven, although information communicated to the relevant individuals, and any consent they have given, will again be factors.

21. **Societal harm** can arise directly from business activity. But it is more likely where the personal information, quite possibly obtained legally or otherwise from businesses, is used by governmental bodies. It includes:

- damage to democratic institutions, for example excessive state or police power; and
- loss of social trust (“who knows what about whom?”).

V. A Matrix to Link Threats and Harms

22. Risk assessment and risk management call for judgment, based upon honest, well-informed and justifiable answers to structured questions about threats and harms. A framework is needed to identify, link and prioritise the various types of threat and harm, ideally in a way that can be easily understood by large and small businesses, by public bodies, by regulators and by individuals.
23. The two draft matrices suggested in Annex 1 demonstrate possible ways of how this might be accomplished in practice. They have been designed as a way of putting privacy on corporate risk radars and getting organisations at least to think about the impact of their activities on the individuals with whom they deal and on the wider community. A framework on these premises—using a common referential—could be initially tested in different contexts by different organisations, not least reflecting varying levels of sophistication and risk aversion. The framework might then mature into a standard template that may, in due course, receive some form of regulatory endorsement to signal a commonly agreed upon approach and become attractive for both organisations and regulators.

VI. A Matrix as an Organisational Risk Management Tool

24. It is envisaged that as a new service, product, technology or activity is developed, a business could use a matrix along the lines suggested in Annex 1 to raise questions and structure a series of judgments arising from each of its inter-sections. Each inter-section requires two specific judgments to be made. A numerical scale would add calibration and rigor:
 - i. How **likely** could this harm arise from any relevant threat? Can this be sensibly quantified on a numeric scale?
 - ii. How **serious** would this harm be if it arose from the threat? Can this be sensibly quantified on a numeric scale?
25. Both judgments should be informed by as much hard data and evidence as possible, such as the nature and volume of the data, consumer complaints, consumer perception research or survey results, industry norms, etc. Also, regulatory guidance could provide an important source of relevant data and regulatory expectations relating to likelihood and seriousness of particular harms, including those affecting fundamental rights and freedoms. The point has already been made that both assessments must be applied objectively, using the reasonable person test. Tangible damage will be objective and usually easier to assess but, even for intangible distress, assessments cannot be based on subjective perceptions. Both the likelihood and the seriousness judgments, however, can and should reflect—and feed into the equation—a prospect of serious harm to a few individuals or less significant harm to many individuals.
26. Each intersection—How likely? How serious?—is a function of the level of the threat and the likelihood that the threat will cause harm. The key judgment is whether there is a **significant risk**. In other words, is there ***a significant likelihood that the particular threat could lead to the particular harm with a significant***

degree of seriousness? Different businesses will have different degrees of risk aversion. Subject to any guidance from its regulator (see below), each business will wish to decide where to fix the level at which a risk is judged to be significant.

27. Where the judgment is made that there is a significant risk—typically as part of an on-going process of risk assessment—appropriate action is then needed to mitigate the risk and implement safeguards to protect individuals from these risks. This might, for example, involve a change of scope, specific safeguards for individuals, or the adoption of a new or improved comprehensive privacy program. A further, post-mitigation, assessment would then be required.

VII. Factoring in the Benefits

28. Not only are some threats to privacy more serious than others, privacy itself is not an absolute value, nor is it the only fundamental right. It must be balanced against other human rights, such as personal security and freedom of expression. There is also a need to strike the right balance with the benefits that arise from the public and commercial uses of personal information. The benefits may flow directly to the individuals concerned or they may accrue at a more societal level, e.g. medical research, law enforcement or improved living standards.
29. As part of the process of assessing the nature and extent of privacy risks, it is necessary to factor in the corresponding benefits because understanding the benefits can help to mitigate risks. Risk cannot be eliminated entirely; and even where it is judged that significant risks exist or remain, there will be situations where the benefits sufficiently outweigh the risks.
30. Benefits may accrue to an individual, to the relevant group of individuals, or to a wider public value and society. Benefits to the business alone are unlikely to outweigh a significant risk to individuals, unless those risks are mitigated and specific safeguards implemented. The important point is that the specific benefits must be:
- identified;
 - articulated;
 - justified by reference to the appropriate external criteria; and
 - judged to outweigh the risk.
31. The accountable business must stand ready to demonstrate how that judgment was reached, producing, as appropriate, the relevant information and evidence upon which it relied.

VIII. The Matrix as a Tool to Prioritise and Guide Regulatory Intervention

32. Though regulators cannot do everything, their responsibilities and challenges are growing while their resources are limited and sometimes in decline. They must be

Selective to be Effective. They need to concentrate on the serious, not the trivial. How should they set their priorities? How, in particular, do they decide which businesses or activities to target for preventative or enforcement intervention?

33. Here is where a risk management matrix may be useful as a tool for regulators. A consensus based on the language and methodology of a matrix could help regulators fix and communicate their own priorities for interpreting and enforcing the rules. This would be welcomed by businesses as it would give them improved predictability and a better idea of where to focus their own risk assessments. At a minimum, the businesses could adopt a “mirror-image” approach in their efforts to avoid regulatory trouble and exceed compliance requirements.
34. One could speculate on various possibilities for a regulator which adopts or endorses a matrix as its starting point:
 - The regulator could signal that it will use that matrix to target industry sectors, particular businesses or activities—anticipating action where it concludes that there is a significant likelihood that a particular threat could lead to a particular harm with a significant degree of seriousness.
 - The regulator could indicate that it expects, as a matter of due diligence, all or some businesses to conduct a risk management exercise on these lines, concentrating regulatory attention on situations where a satisfactory exercise has *not* been conducted.
 - The regulator could use the matrix to determine whether the business has adopted appropriate risk mitigations (e.g. limitations, restrictions, safeguards).
35. There would be a further advantage if regulators could communicate tolerance levels to help businesses decide whether a risk is significant. It would be a powerful message, for example, for a regulator to state that, for a particular type of activity, a risk would be significant where the assessment score exceeds a prescribed level.
36. The approach implies that regulators will need to assess the efficacy of risk processes. In keeping with the principle that risks are usually mitigated but seldom eliminated, there may be situations where a regulator concludes the risk assessment process was reasonable and complete but simply disagrees with the end decision. In that situation, the company may be especially exposed if the harm in fact materialises. In other situations, as the FTC has shown with imaginative use of consent decrees which impact on privacy programs, a regulator that finds fault with risk assessments is well placed to ensure constructively that its rulings have positive effects in the future.
37. Finally, a common and consistent use of a risk management matrix by regulators in different countries would lead to much needed consistency and even harmonisation of expected outcomes, even in situations where the underlying rules may not be always the same. The potential for the risk-based approach and risk management matrix to be used globally would be a powerful step forward towards global interoperability.

IX. The Matrix as a Tool Where a Harm Has Been Suffered

38. There is also scope for regulators and courts to use a risk management matrix as a remedial tool where a harm has actually been suffered by one or more individuals. It may help determine how the harm came about, not least in efforts to repeat similar incidents. More directly, the matrix could influence the nature and scale of regulatory sanctions or compensatory redress. It may, in particular, help to decide the foreseeability of the harm that arose from the threat.
39. In a world where regulators rightly pay most attention to those who knowingly ignore their obligations, are cavalier or are repeat offenders, they are entitled to ask for evidence that a risk assessment had been conducted before the activity in question was launched.
40. It is to be expected that a risk-based approach would be accompanied by significantly heavier penalties and sanctions where risk assessment has been nonexistent or manifestly inadequate.

X. Implications for Lawmakers

41. This paper focuses on a risk-based approach as a means of implementing and calibrating existing legal requirements and compliance in practice, to make them more effective. As such, the paper advocates adopting a risk-based approach which may be attractive to businesses and regulators, within the frameworks of current legislation. If this proves to be an effective way of maintaining and improving protections in practice, it might be contemplated in the longer term that suitable legislative text could be developed to embrace more comprehensively a risk-based approach in preference to some more rigid and prescriptive, which may be judged as ineffective. This is not, however, the focus of this paper.

XI. Issues for Further Consideration

42. The above discussion raises a number of issues and questions that require further consideration as a part of the Centre's Privacy Risk Framework Project:
 - Any methodology for risk assessment needs to have an agreed definition of "risk". Do we mean risk to privacy, or risk to personal data protection? Do we mean risk to individuals' other rights and freedoms?
 - Can and should organisations consider societal harms in their risk assessments?
 - Who decides what is "risky" and what is "harm"? When and how do they decide? Are there categories of processing that are considered *per se* risky or that are always considered harmful? Can the potential risks and harms associated with certain data processing be assessed by the controller without inserting too much subjectivity? If so, how do we ensure sufficient legal certainty, both for the organisation and for the individuals?

- What is the role of the affected individual in any risk assessment? How much and what kind of participation or transparency is required?
 - Do individuals need to be told about or allowed to participate in the risk assessment? Should companies share the outcomes of any risk assessment with individuals?
 - Should individuals be given an opportunity to object to the outcomes of a risk assessment? Should they have a right to object to processing despite a contrary risk assessment?
 - What is the role of consent in this context? Does consent trump risk assessment or vice-versa?
- What do we know about individuals' perception of risk and harm to themselves? Are there surveys and market research, and is there sufficient existing knowledge in the business community? What are other ways to obtain relevant information on this topic? Can we monitor the reaction of individuals over time? Can we use social media as a channel for such monitoring?
- We should not replace one bureaucracy with another. The proposed EU Data Protection Regulation aims to reduce bureaucracy. Would an elaborate and documented case-by-case risk assessment for every processing of personal data be impractical as well as cost and labor intensive? On the other hand, given that many privacy laws (including the EU Directive and the proposed EU Regulation) already require risk assessments in many instances—such as in the “legitimate interest” balancing test under the current EU Directive—a widely agreed upon risk assessment methodology may improve efficiency and reduce administrative burdens.
- Will the risk-based privacy framework be scalable for SMEs, who are some of the main drivers of innovation and new technologies and services?
- Companies already routinely assess privacy risks to themselves, such as non-compliance, reputational and litigation risks. How do these types of risk assessments relate to those that focus on risks to individuals, particularly where these may not overlap? How does an organisation integrate both risk assessments seamlessly?
- Risk assessment is an integral part of organisational accountability. How, exactly, will a risk-assessment methodology help demonstrate accountability and compliance with applicable legal requirements?
- What is the role of regulators *vis-à-vis* an organisation's decision to process data based on a risk analysis? The risk-based approach must not undermine a regulator's ability to challenge the validity of risk-analysis outcome. Can risk analysis outcomes be challenged even in the absence of harm?
- What is the role of technology and technologists in developing and implementing risk-based solutions to privacy protection?

- Any risk methodology and framework must be capable of being exported and used by technologists, data scientists, data anthropologists, engineers and many others who, normally, will not have an intuitive or developed understanding of privacy issues. How can we socialise the risk-analysis concept more broadly and work with nonprivacy practitioners to that end?
- Data ethics is a new discipline. How does an ethical decision-making model fit or should be reflected in any risk assessment methodology?

XII. Next Steps

43. The Centre will continue work towards a comprehensive privacy risk framework, drawing on the expertise of its members, project participants and privacy experts, including from academia and the regulators community. It will also seek to collaborate with other organisations interested in the risk-based approach to privacy.

44. Future work on the project may include:

- developing additional discussion papers based on further study of all issues identified in this paper or raised by the risk-based approach to privacy;
- holding further workshops to receive input and discuss our learnings;
- examining existing risk analysis practices to inform the development of the privacy risk framework;
- taking stock of current laws and regulatory schemes that require and incorporate risk analysis today;
- identifying new areas of potential use for the risk-based approach, such as in response to new and evolving privacy threats in the modern data economy;
- refining and developing the practical tools associated with risk analysis, such as the risk matrix, and thinking about practical implementation of the risk-based approach;
- undertaking case studies on applying the risk methodology under development in the present project, including the risk management matrix, to various real-life scenarios, such as in activities involving: health data, big data, anonymised/pseudonymised data; new products and services, etc.;
- considering individual participation and transparency issues;
- examining the potential uses of the risk-based approach by the different privacy stakeholders—organisations, regulators, and law and policy makers; and
- studying the potential of the risk-based approach to enable global interoperability.

Version 1.0 06/2014										
DRAFT - Risk Matrix										
Risks	Unjustifiable Collection			Inappropriate Use			Security Breach			Aggregate
				Inaccuracies Not expected by individual Viewed as Unreasonable Viewed as Unjustified			Lost Data Stolen Data Access Violation			
	Likely	Serious	Score	Likely	Serious	Score	Likely	Serious	Score	Risk Rank
<u>Tangible Harm</u>										
Bodily Harm	0	0	0	0	0	0	0	0	0	0
Loss of liberty or freedom	0	0	0	0	0	0	0	0	0	0
Financial loss	0	0	0	0	0	0	0	0	0	0
Other tangible loss	0	0	0	0	0	0	0	0	0	0
<u>Intangible Distress</u>										
Excessive surveillance	0	0	0	0	0	0	0	0	0	0
Suppress free speech	0	0	0	0	0	0	0	0	0	0
Suppress associations	0	0	0	0	0	0	0	0	0	0
Embarrassment/anxiety	0	0	0	0	0	0	0	0	0	0
Discrimination	0	0	0	0	0	0	0	0	0	0
Excessive state power	0	0	0	0	0	0	0	0	0	0
Loss of social trust	0	0	0	0	0	0	0	0	0	0

Legend:

Rank 'Likely' from 10 (high) to 1 (low) based on the highest score for any component

Rank 'Serious' from 10 (high) to 1 (low) based on the highest score for any component

Aggregate Risk Rank:

Highest score is 300

Lowest score is 0

Draft Risk Matrix

Proposed Processing:		THREATS													
		Unjustifiable Collection of Data	Inappropriate Use of Data										In Wrong Hands		
			Storage or use of inaccurate or outdated data	Use of data beyond individuals' reasonable expectations	Unusual use of data beyond societal norms, where any reasonable individual in this context would object	Unjustifiable inference or decision-making, that the organisation cannot objectively defend	Lost or stolen data	Data that is unjustifiably accessed, transferred, shared or published							
HARMS	Tangible Harm														
	Bodily harm	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	
	Loss of liberty or freedom of movement	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	

Draft Risk Matrix

	Damage to earning power	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	
	Other significant damage to economic interests	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	
	Intangible Distress												
	Detriment arising from monitoring or exposure of identity, characteristics, activity, associations or opinions	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	
	Chilling effect on freedom of speech, association, etc.	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	

Draft Risk Matrix

	Reputational harm	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	
	Personal, family, workplace or social fear, embarrassment or anxiety	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	
	Unacceptable intrusion into private life	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	
	Discrimination or stigmatisation	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	

Draft Risk Matrix

Societal Harm															
	Damage to democratic institutions (e.g. excessive state or police power)	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	
	Loss of social trust (Who knows what about whom?)	how likely?		how likely?		how likely?		how likely?		how likely?		how likely?		how likely?	
		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?		how serious?	

Annex 2: Examples (non-exhaustive) of Risk Assessment Schemes Used by UK Regulators

- Health and Safety Executive / Local Authorities Enforcement Liaison Committee Priority Planning system;
- Office of Fair Trading—Trading Standards Risk Assessment Scheme;
- Food Standards Agency—Food Hygiene and Food Standards Intervention Rating Schemes; and
- Local Authority Integrated Pollution Prevention and Control (LA-IPPC) and Local Authority Pollution Prevention and Control (LAPPC) Risk Methodology.