IAF

Information Accountability Foundation

RE: INVITATION FOR PRELIMINARY COMMENTS ON PROPOSED RULEMAKING CYBERSECURITY AUDITS, RISK ASSESSMENTS, AND AUTOMATED DECISIONMAKING For the California Privacy Protection Agency, PR 02-2023

The Information Accountability Foundation (IAF) is a non-profit research and educational organization headquartered in Los Gatos, California. It was created in 2013 to encourage fair information usage so that data pertaining to people might create real value for those people in a protective manner. The IAF is the incorporation of the Global Accountability Dialog that created "The Essential Elements of Accountability" that have been codified in the EU General Data Protection Regulation (GDPR), Colombia and Mexico privacy laws, and guidance in numerous countries. Accountability requires organizations to be responsible and answerable for their data use.

Assessments are central to organizations using data responsibly. Conducting assessments also create the record that organizations are accountable. To build accountability into advanced analytics, the IAF authored the "The Unified Ethical Frame for Big Data Analytics" that placed burdens on data users to assess the risk those organizations created for others. Since 2014, the IAF has worked with stakeholders to create assessment templates in the United States, Europe, Hong Kong, and Canada. The IAF work has inspired assessments in other jurisdictions as well. Appendix Part B includes links to many of those assessment templates and Part C on Enforcement of assessments. The IAF currently is working on assessments that look to the full range of interests required by the final privacy rules just issued in Colorado. It is from that nine years' experience in developing assessments in collaboration with the full range of stakeholders that the IAF provides comments.

The IAF focuses its comments on Section II and III. The IAF uses the questions of the California Privacy Protection Agency (CPPA) as the starting point for the IAF's answers.

II. RISK ASSESSMENTS

1. What laws or other requirements that currently apply to businesses or organizations (individually or as members of specific sectors) processing consumer' personal information require risk assessments?

   Risk assessments related to the use of data pertaining to people come in many forms. There are privacy impact assessments (PIAs), data protection impact assessments (DPIAs), ethical assessments, legitimate interest assessments, and increasingly algorithmic assessments. PIAs were suggested strongly in 2012 by Canadian regulators in "Getting Accountability Right Through a Comprehensive Privacy Management Program". This document inspired similar documents in Hong Kong and Colombia. Many large Canadian organizations adopted PIAs in response to this regulatory encouragement. It was not a legal requirement. Those PIAs focused on data subject rights and today fall short of what seems to be required in the California law.

   The GDPR requires legitimate interest assessments that balance the legitimate interests of the controller against the full range of rights and interests of the data subject. That requirement has had mixed success both in governing legitimate interests as a successful legal basis and in bringing the full range of stakeholders into consideration.

The GDPR also requires DPIAs when a process creates “high risk” for data subjects. As referenced in the request for comments, the European Data Protection Board (EDPB) has published guidance on when and how to do DPIAs. While the EDPB guidance differentiates between risk to the organization, that is in part enterprise risk management, and risk to data subjects, the EDPB guidance doesn’t define what risk means. Given that gap, the IAF conducted a project called “Risk of What?” Are regulators looking for impediments to exercising data subject rights, such as transparency and data minimization, or inappropriate bad outcomes to people, as is the basis for the U.S. state and federal Fair Credit Reporting Acts (FCRAs)? Whatever the experience in Europe has been, it is inadequate because European regulators have not embraced totally Recital 4 of the GDPR which requires consideration of all stakeholders and the balancing of all fundamental rights.

In recent weeks, the Colorado Attorney General adopted final rules pursuant to the Colorado Privacy Act. Rule 8.04 provides guidance on Data Protection Assessment Content. Number 6 under that rule defines sources and nature of risks to the rights of consumers. That section seems to reflect the “Catalog of Problematic Data Actions and Problems” contained in the “NIST Privacy Framework: A Tool for Improving Privacy Through Enterprise Risk Management.” The IAF believes the NIST catalog is an excellent place to start when defining the risks to people and society when data pertaining to people is processed. The IAF used that catalog when developing its list of “Adverse Processing Impacts and Defining Risk” as part of the IAF model legislation, the FAIR ANF OPEN USE ACT. The Agency may also find the NIST CPPA-CPRA Crosswalk helpful.

The chart below cross references the risks identified in the Colorado Rules against the IAF Adverse Processing Impacts.

Colorado Privacy Act Harms mapped to IAF Adverse Processing Impacts

Colorado Privacy Act Rules PART 8.04(6) – Privacy Harms	**IAF-defined Adverse Processing Impacts** *(Derived from NIST Catalog of Problematic Data Actions)*
a. Constitutional harms, such as speech harms or associational harms;	(9) Loss of autonomy and (10) Other detrimental or negative consequences
b. Intellectual privacy harms, such as the creation of negative inferences about an individual based on what an individual reads, learns, or debates;	(6) Stigmatization - Stigmatization or reputational injury
c. Data security harms, such as unauthorized access or adversarial use;	*Security breaches may cause outcomes from harmful processing that may take place when a breach occurs but are not a direct harm to individuals. Adequate security requirements should be covered elsewhere in a regulation.*
d. Discrimination harms, such as a violation of federal antidiscrimination laws or antidiscrimination laws of any state or political subdivision thereof, or unlawful disparate impact;	(8) Discrimination - Discrimination in violation of Federal antidiscrimination laws or in laws of any State
e. Unfair, unconscionable, or deceptive treatment;	Includes all adverse processing impacts including (4) Inconvenience or expenditure of time

f. A negative outcome or decision with respect to an individual's eligibility for a right, privilege, or benefit related to financial or lending services, housing, insurance, education enrollment or opportunity, criminal justice, employment opportunities, health-care services, or access to essential goods or services;	(5) A negative outcomes or decision with respect to an individual's eligibility for a right, privilege or benefit – Denial of employment, credit, insurance, a license, etc.
g. Financial injury or economic harm;	(1) Financial Loss - Direct or indirect financial loss or economic harm
h. Physical injury, harassment, or threat to an individual or property;	(2) Physical Harm - Physical harm, harassment, or threat to an individual or property
i. Privacy harms, such as physical or other intrusion upon the solitude or seclusion or the private affairs or concerns of Consumers, stigmatization or reputational injury;	(6) Stigmatization - Stigmatization or reputational injury (9) Loss of Autonomy - Loss of autonomy through acts or practices that are not reasonably foreseeable
j. Psychological harm, including anxiety, embarrassment, fear, and other mental trauma; or	(3) Psychological Harm - Psychological harm, including anxiety, embarrassment fear, and other mental trauma
k. Other detrimental or negative consequences that affect an individual's private life, private affairs, private family matters or similar concerns, including actions and communications within an individual's home or similar physical, online, or digital location, where an individual has a reasonable expectation that Personal Data or other data will not be collected, observed, or used.	(10) Other detrimental or negative consequences

The IAF believes the CPPA should begin its regulations on risk assessments with a regulation similar to the rule enacted by the Colorado Attorney General. Having a set of common risks would enhance the ability for organizations of all sizes to get it right when trying to determine if a processing is highly risky.

The Colorado rule requires organizations to assess the risk to the individual to whom the data pertains, risk to groups of individuals, and risk to society as a whole. The business community has limited experience in looking beyond the risk to the business and the risk to data subjects. This comprehensive approach will require assessments begin with clearly thinking through and articulating the relevant stakeholders, how they might be impacted, and to what effect. The IAF believes that the Colorado rules will create the encouragement for that type of assessment to develop. As mentioned earlier, the IAF has developed templates for these types of assessments in the past.

Lastly, the EDPB guidance references the fact that organizations need to understand how to review their activities to determine whether a DPIA is necessary. The IAF believe that type of guidance would be useful as part of the regulations the CPPA issues.

2. What harms, if any, are particular individuals or communities likely to experience from a business's processing of personal information? What processing of personal information is likely to be harmful to these individuals or communities, and why?

The chart that is part of the answer to question 1 lists adverse processing impacts. What is missing from Question 2 is the harm to people of not processing information. Organizations make decisions every day to not process information pertaining to people because of compliance concerns related to secondary use of data. Rules should be balanced to look at both sides of the risk equation. Are the data pertaining to people that do not get processed because they are a secondary use more or less harmful to society? Instead of a flat prohibition on secondary use, that kind of balancing should be done.

3. To determine what processing of personal information presents significant risk to consumers' privacy or security.

a. What would be the benefits and drawbacks be of the Agency following the approach outlined in the EDPB's Guideline on Data Protection Impact Assessments.

As discussed above, the EDPB guidance only looks at the risk to the data subject, not the risk to other stakeholders. Also, the EDPB guidance does not catalog the risks that might come from the processing of data or not processing data. This balancing is important increasingly when determining the productive use of AI, the quality of complete data sets, and the concerns about profiling.

b. What other models or factors should the Agency consider? Why? How?

The IAF suggests the CPPA consider Colorado Rule 8.04 that includes assessing the full range of stakeholders for the risk factors described in the rule.

c. Should the models or factors be different, or assessed differently, for determining when processing requires a risk assessment versus a cybersecurity audit? Why, or why not? If so, how?

The IAF is not addressing cybersecurity issues.

d. What processing, if any, does not present significant risk to consumer's privacy or security? Why?

Every time data pertaining to a person is used there is risk. Organizations should triage a processing to determine the level of risk both to the protection of the data and the protection of the people to whom the data pertains. It is the context for the use that ultimately defines the risk level. Whitelists and blacklists have limited utility in a fast-evolving world.

4. What minimum content should be required in business's risk assessments? In addition:

a. What would the benefits and drawbacks be if the Agency considered the data protection impact assessment content requirements under the GDPR and the Colorado Privacy Act?

b. What, if any, additional content should be included in risk assessments for processing that involves automated decision making, including profiling? Why?

The IAF already has suggested that Colorado Rule 8.04 is a good place for the CPPA to start it rulemaking. The IAF is developing an assessment template for Colorado assessments that is not ready for this submission. The IAF also is developing the concept of assessing on the three dimensions of stakeholders, their fundamental interests, and adverse consequences to those fundamental interests.

Probabilistics, the basic process behind profiling, has been accelerating since the development of the first bankruptcy scores in the late 1980's. Automated decision-making is a natural development of quickly getting to decisions where probabilities are clear. However, the fact that an outcome is probable is different than it being certain. The federal and state FCRAs have done a very good job of describing where decisions have a legal or similarly significant effect. Probabilistics add questions to the assessment process. There should be continuity from a base assessment to anything that needs to be added for profiling and automated decision-making.

5. What would the benefits and drawbacks be for businesses and consumers if the Agency accepted businesses' submission of risk assessments that were complete in compliance with GDPR's or Colorado Privacy Act's requirements for these assessments? How would businesses demonstrate to the Agency that these assessments comply with CCPA's requirements?

   Organizations already looking at where the GDPR and the Colorado requirements overlap and where they differ. It is impractical for organizations to conduct different assessments for the EU and Colorado (and California in the future). A Colorado assessment may begin with the GDPR factors and add the requirements related to the full range of stakeholders and adverse consequences in Colorado Rule 8.04. It would do the same thing with any additional California requirements. Fundamentally, the GDPR, the Colorado rules and the CPRA all call for the same thing: the conduct of assessments that consider whether risky processing is being conducted (risky processing includes the processing of sensitive personal data), evaluation of the benefits versus the risks of processing personal data for the business, its consumers, the public, and other stakeholder, and the avoidance of processing activities if they place significant potential risks on data privacy, outweighing its overall benefits.

   The CPPA then should spot check assessments to make a judgement whether they were developed competently and with integrity.

6. In what format should businesses submit risk assessment to the Agency? In particular:
   a. If business were required to submit a summary risk assessment to the Agency on a regular basis (as an alternative to submitting every risk assessment conducted by the business.

      If organizations are required to submit every risk assessment to the CPPA, the CPPA will be flooded with submissions and will have limited ability to review those submissions. Informal conversations with organizations have led the IAF to believe that European agencies receive very few DPIAs because they must be submitted only if there is significant residual risk. Once risks are identified, organizations typically modify processing to mitigate those "significant risks."

      The development of a summary risk assessment format should be a separate regulatory undertaking by the CPPA. Some jurisdictions are thinking about using the code of conduct process as a means for establishing the content of a summary assessment.

7. Should compliance requirements for risk assessments or cybersecurity audits be different for businesses that have less than $25 million in annual gross revenues? If so, why, and how?

   The IAF is not responding to this question.

8. What else should the Agency consider in drafting its regulations for risk assessments.

   Organizations must develop a continuous process for determining the level of risk they create for others when processing data. There are discussions concurrently about fair AI assessments, ethical assessments, and algorithmic assessments. Risk assessments should be part of a seamless process that begins with triage on whether a processing is going to create risks of adverse processing for identified stakeholders.

III. AUTOMATED DECISIONMAKING

The February 10, 2023, Invitation for Preliminary Comments asks a series of questions related to automated decision-making and profiling. The IAF is not responding to the specific questions but instead setting forth some basics for the discussion. The fact is that automated decision-making is baked into how things work on an everyday basis. For example, the CPPA uses automated decision-making on requests from browsers to access the CPPA's servers on a daily basis. These decisions have the effect of limiting who can browse the CPPA's website and file complaints. This is good because the alternative would be constant security breaches. However, the issues related to profiling and automated decision-making predate when consumer browsers made the Internet a consumer medium.

Martin Abrams, former Founder and President, currently the Chief Policy Innovation Officer of the IAF, was the President of the Centre for Information Policy Leadership (CIPL), the Vice President of Experian Policy Solutions, and the Assistant Vice President and Community Affairs Officer of the Cleveland Federal Reserve Bank. His background gives him the perspective to provide the following comments.

The consumer Internet accelerated an observational age that in turn accelerated the use of data for probabilistics pertaining to how people behave. The first broad-based probabilistic use of consumer data was probably the Fair Isaac credit risk score in 1989. It was quickly adopted by the consumer lending industry as an aid to better decisioning than was possible with the subjectivity of decisions made purely by lending officers. Soon that aid to people evolved into automated credit decisions. The U.S. Department of Justice (DOJ) investigated whether those decisions had the effect of making decisions on grounds that violated the Equal Credit Opportunity Act (ECOA). Since the data for credit risk scores came directly from credit bureaus, the FCRA required that the use of scores must be disclosed along with the factors that led to the denial. So, from the very beginning, the use of profiling and automated decision-making for substantive decisions were covered by a fair processing law, the FCRA.

In Europe, there was no uniformity in the data available for consumer credit decision-making. As Europe evolved towards the creation of the 1995 EU Privacy Directive, there were debates on whether it was unseemly for decisions on people to be made solely by a machine. Those concepts on what is seemly or not influenced the drafting of Article 22 of the GDPR. So, there are cultural differences between the way that Europe sees these issues and the way they are seen in the United States. The fact is that the relationship between profiling, the use of probabilistics against broad data sets, and automated decision-making is muddled still under Article 22 of the GDPR.

The 21st century saw the rise of analytic skills that allowed for the use of unstructured data into advanced analytic processes. Legacy statistics tested causality, while the growth of big data switched the dominant theme to correlation. This change naturally raised questions about the accuracy of the correlations, whether they were appropriate to apply, and whether they were influenced by the bias built into available data sets. This development has informed the debate about algorithmic fairness. These concerns have accelerated with the growing use of AI, which is the next stage of advanced analytics in our observational world.

So, in thinking about the questions the CPPA is asking, some pragmatic truths need to be addressed:

- Profiling is probabilistics built with consumer data. Building choice into the data that feeds the probabilistics has the unintended consequences of skewing the accuracy of predictive values. Choice worked when the relationship was one on one. Most relationships are no longer one on one. Ours is an observational world where there are not many one-on-one relationships. Choice no longer fits and indeed harms the process in an observational world.
- Automated decision-making is built into how many modern processes work, including the functioning of the CPPA's cybersecurity processes. Many automated decision-making processes are subject already to laws such as the FCRA, ECOA, and Fair Housing Act (FHA). The FCRA, ECOA, and FHA wrestled with these issues already and decided that the benefits of the automated decision-making outweighed the risks. Those Acts have methods for determining whether the automated decision-making is biased or not (after the fact testing), and those methods are just as applicable today as they were when they were implemented.
- Much of the emotions that pertain to automated decision-making are related directly to whether one thinks it is fairer for a person to make a decision or whether a well-governed algorithm, in the end, would be fairer. As mentioned above, the DOJ in the context of the ECOA decided that a well-governed algorithm was better.

The IAF staff believes this is where the discussion should begin. Thank you for the opportunity to contribute comments to this important rulemaking process.

Sincerely,

Martin Abrams, IAF Chief Policy Innovation Officer
Barbara Lawler, IAF President
Lynn Goldstein, IAF Senior Strategist

March 27, 2023