ROBERT GELLMAN
Privacy and Information Policy Consultant **202-543-7923**
419 Fifth Street SE **bob@bobgellman.com**
Washington, DC 20003 **www.bobgellman.com**

Three Bad Ideas in the Consumer Privacy Bill of Rights

March 5, 2015
Version 1.4

The Obama Administration released its Consumer Privacy Bill of Rights (CPBR) on a Friday afternoon (2/27/15). In Washington, the late Friday press release is a classic way of trying to avoid attention and comment. The privacy community in general did not like the bill, with some calling it a step backward from existing consumer privacy protections. Some politely welcomed the bill as a vehicle for further discussions, but no privacy or consumer group showed the slightest support for the actual proposal. Parts of the business community criticized the bill, but I suspect that those paying attention secretly wish it could pass and preclude a stronger bill. No one in the EU will be fooled into thinking that the CPBR comes anywhere close to meeting European adequacy standards.

It's hard to tell if the CPBR represents a serious effort because it strikes me as more of a *privacy-prevention* law. I say that because it mostly proposes privacy controls that range from weak to non-existent and then for the most part preempts better state laws. I will limit myself here to three main issues raised by the proposal: the multistakeholder process, the idea of "context" and privacy risk management.

1. Multistakeholder process. The bill seeks to enshrine in law the multistakeholder process that the National Telecommunications and Information Administration (NTIA, part of the Department of Commerce) established in 2012. The originally announced goal of the multistakeholder process was "open, transparent forums in which stakeholders who share an interest in specific markets or business contexts will work toward consensus on appropriate, legally enforceable codes of conduct." The topic of the first effort was mobile application transparency. An ongoing effort addresses commercial use of facial recognition technology.

If you assumed that the legislative proposal resulted from some demonstrated success of the multistakeholder process, you were wrong. The first NTIA multistakeholder process developed a transparency code for mobile apps, a narrow subset of fair information practices. Susan Grant from Consumer Federation of America best described the shortcomings of the process in her comments about the mobile app transparency code:

> It is not surprising that the product is so flawed given the problems with the process itself. There was never any clear procedure for how it would work and what would constitute success. There was no legal framework on which the code could be built, so that even terms such as "user data" are not clear and universally understood. The last meeting of the stakeholder group yesterday was as confusing as the process has been all along, with a "vote" being taken that allowed multiple attendees from the same companies or organizations to vote and

> resulted in no clear consensus. The groups that drafted the code, a small subset of the stakeholders, simply declared victory and the process ended.

The legislation solves none of the demonstrated weaknesses of the multistakeholder process. There are at least three big problems. First, there is no formal procedure for adoption of a code. It wasn't clear during the first multistakeholder effort (which dragged on for more than a year) when the "code" would be ripe for a vote. NTIA just pushed things along, and a vote occurred even through there was no apparent consensus. In other words, NTIA declared victory and moved on.

Second, anyone can participate, and anyone can vote. While there's nothing wrong with broad participation, the lack of rules governing representation (who represents what interest) is a real problem. Consumer and privacy groups do not have the resources to participate in multiple multistakeholder processes. Companies can send staff, hire lawyers to represent them, or rely on trade associations. Industry can send as many people to meetings as it chooses, and it can easily overwhelm any other participants by sheer numbers. As a result, the multistakeholder playing field seems inherently unequal, and the lack of procedures only makes this worse. The more processes that occur, the more unequal it is for the public interest community.

Third, the Commerce Department is not a neutral party. It functions as a representative of the business community in all of its activities, and it puts its thumb on the multistakeholder scale by dictating procedures and deciding agendas. No equivalent Administration office weighs in on behalf of consumers. The business bias of the Department of Commerce is evident in the CPBR's narrowly circumscribed protections for consumers.

Can a multistakeholder process work under other circumstances? Maybe, but it is hard. I observe that a somewhat similar endeavor sponsored by the World Wide Web Consortium to develop standards for do-not-track browser controls has dragged for years with no apparent hope of a consensus result. No process will produce a positive result unless there is a realistic threat of regulation to force business to compromise. Also, the result of the process has to cover everyone, not just those who volunteer to comply. There's no point in a process that results in applying privacy rules to ten percent of an affected industry.

2. Context. The CPBR places a lot of weight on the notion of *context*. Professor Helen Nissenbaum of New York University first put forward context as a way of explaining why consumers express privacy concerns differently and inconsistently. An easy example is that an individual might tell a physician something that she would not disclose to her spouse. In the context of the practice of medicine, people have a sense of expectation about what happens to their information and how it will be protected. (That few people actually understand how widely health information flows inside and outside the health care system is a point I cannot pursue here.) The lack of consumer understanding of industry data practices makes it hard to rely on context for writing standards.

So while context is useful as a way of explaining consumer behavior, context does not work as a legislative standard for controlling the processing of personal information. The CPBR provides a complex definition for context that is worth reproducing here:

> (k) "Context" means the circumstances surrounding a covered entity's processing of personal data, including but not limited to—
>
> (1) the extent and frequency of direct interactions between individuals and the covered entity, if any;
> (2) the nature and history of the interactions described in paragraph (1);
> (3) the level of understanding that reasonable users of the covered entity's goods or services would have of how the covered entity processes the personal data that it collects, including through any notice provided by the covered entity;
> (4) the range of goods or services that the covered entity offers, the use of such goods or services by individuals, the benefits of such goods or services to individuals, and the brand names that the covered entity uses to offer such goods or services;
> (5) information known by the covered entity about the privacy preferences of individual users of its goods or services;
> (6) the types of personal data foreseeably processed in order to provide a good or service that an individual requests from the covered entity;
> (7) the types of personal data foreseeably processed in order to improve or market a good or service that an individual requests from the covered entity;
> (8) the types of personal data foreseeably processed as customary business records;
> (9) the age and sophistication of individuals who use the covered entity's goods or services, including whether the covered entity's goods or services are directed toward minors or the elderly;
> (10) the extent to which personal data under the control of the covered entity are exposed to public view; and
> (11) the extent to which personal data under the control of the covered entity are obscured.

What's wrong with context as a legislative standard? To begin, this definition has eleven elements, each element is vague, and it is unclear how to weight or apply each element. Different elements suggest contradictory results in any given context. I won't pause here to critique the list further. The bigger problem is that it isn't apparent who will be the arbiter of what is allowable in any given context. It is hard to find any candidate to make decisions other than the covered entity processing the data.

Second, contexts have few clear borders. Imagine that you visit your physician and pay your health insurance co-payment with a debit card. The same transaction information now exists in at least four different contexts. There is a health record, an insurance record, a bank record, and an electronic funds transfer record. Will any consumer understand these contexts, appreciate how different players process personal information in different contexts, or be aware that consumer rights may vary with the context? Compare this with the EU approach where privacy rights are similar everywhere that data protection rules apply.

Third, in many if not most cases, the CPBR covered entity processing the personal information decides unilaterally on the context. This is certain to be true online. If you visit a website and engage in an activity that involves disclosure of personal information, the policy of the website determines the context. The consumer had to "agree" to the policy as a condition of using the website. All the other listed contextual elements are nice, but consumer acceptance trumps them all. The bill says so expressly when it provides that "personal data processing that fulfills an individual's request shall be presumed to be reasonable in light of context." In other words, abandon all rights under the CPBR ye who click here because the covered entity that processes the data defines the context. Existing privacy policies give consumers few firm rights, and the bill may actually even weaken those rights.

Fourth, existing contexts for much personal information processing are significantly weighted against consumers. Internet advertising companies and others who monitor web usage are unknown to consumers so the context of their activities does not have any pretense of consumer involvement. The data broker industry for the most part operates without any effective notice to consumers. Data brokers collect, process, and sell all the consumer information they can obtain without any agreement from consumers or legislative restriction. That is the context in which data brokers have operated for decades. The established context gives brokers all the benefits and all the value, and consumers obtain little, if anything, in return. The legislation will not change data broker behavior or responsibility in any meaningful way. Consumers will continue to have no rights, and the CPBR would cement that lack of rights into law.

By contrast, the Fair Credit Reporting Act, a long and detailed law that limits uses and disclosures of credit reports and defines consumer rights, firmly establishes the context of credit reporting. It took Congress more than 40 years from original passage and several major amendments to strike a decent balance between consumers and credit bureaus with respect to credit reports. A few vague phrases about context are not a substitute for defined rights and clear limits.

3. Privacy Risk Management. The bill introduces the notion of privacy risk management. However, the requirement to assess privacy risks does not apply at all to any processing that is "reasonable" in light of context. Marrying two vague concepts – context and reasonableness – is not likely to produce clear, consistent results that benefit consumers in any way. Virtually no covered entities will have to engage in privacy risk management because they determine both the context and whether their activities are reasonable. Privacy risk management only applies to conduct that is *unreasonable* in light of context.

The required privacy risk analysis lists several elements that covered entities must examine for privacy risk. Covered entities must mitigate any identified privacy risks through heightened transparency, individual control, and perhaps more.

If that sounds okay, wait. It gets worse. A covered entity does not have to provide more transparency or individual control or otherwise mitigate privacy risks if it can get a Privacy Review Board to sign off on the activity and agree broadly that the benefits outweigh the risks. That is, the benefits to the covered entity processing the data outweigh the risks to data subjects.

How can the covered entity possibly get approval from a Privacy Review Board for conduct that is unreasonable in light of context? The answer is simple. The covered entity appoints and pays the Privacy Review Board. There is no requirement for public disclosure of a Privacy Review Board's activities, decisions, or conflicts of interest. There are no established requirements or principles to guide the Board's work. There is no requirement for participation by consumers, although a token consumer representative would make no difference anyway. The covered entity that created the board would have total control of the process.

In other words, the privacy risk management process is the privacy equivalent of a show trial. The covered entity can be the prosecutor, judge, and jury. The board can function without any public oversight. Outcomes will be preordained.

The model for the Privacy Review Board is reportedly the institutional review board (IRB) used to review research proposals. IRBs operate under standards that developed over decades through participation by researchers, ethicist, lawyers, and policy makers. Those who review research proposals are separate and independent from those who want to carry out the research. Nevertheless, IRBs have shortcomings well known in the research world. Still, universities run most IRBs, and universities generally try to do things reasonably well because there is real oversight and enforcement of the IRB process. The research world worries about its regulators at the Office of Human Research Protections because OHRP can (and has!) cut off federal funding for research. There is real tension in the system, which helps it operate in better balance.

Where are the equivalent standards and oversight for Privacy Review Boards? If anyone thinks that the Federal Trade Commission will have the resources to monitor Privacy Review Boards in addition to its existing work, to reviewing and approving numerous multistakeholder processes, and to carrying out other new (and largely meaningless) requirements under the CPBR, I have a bridge in Brooklyn that I would like to sell you.

There is a lot more in the CPBR to dislike, but I promised to limit myself to three issues. A full analysis would exhaust your willingness to read it. The CPBR is largely a privacy shell game that offers only an appearance of meaningful consumer privacy protections.

The best things about CPBR are that no one thinks that it has any meaningful chance of passage and that the entire idea will sink without a trace at the end of the Obama Administration in less than two years.

*****