

## Quantitative Privacy Risk Analysis

R. Jason Cronk  
*Enterprivacy Consulting Group,*  
*FL, USA*  
*rjc@enterprivacy.com*

Stuart S. Shapiro  
*MITRE Corporation,*  
*MA, USA*  
*sshapiro@mitre.org*

**Abstract**—Most privacy risk assessment methodologies are homegrown and qualitative. Numerical models generally involve largely arbitrary quantifications. FAIR, a quantitative risk model for information security related risks, can be modified for privacy, providing more meaningful measurements and supporting comparison of risks of similar scenarios with varying controls to organizational tolerances.

**Index Terms**—privacy risk, FAIR, quantitative risk, risks assessment

### 1. Introduction

What is risk? At an abstract level, it is a function of threats, vulnerabilities those threats can exploit, and the resulting adverse consequences. More fundamentally, it is an effort to grapple with the uncertainties surrounding the potential occurrence and impacts of particular events. At a practical level, risk is an adverse consequence associated with some indicator of the magnitude of severity and a frequency of occurrence. If magnitude and frequency can be quantified, the risk can be nominally quantified as frequency times magnitude. To the extent these quantities reflect some objective grounding, statistical concepts become applicable. If either of these cannot be quantified, qualitative mappings are typically employed.

Privacy risks arise from interactions with individuals or their proxies. Proxies can include, but are not necessarily limited to, information about that individual, an individual's property or their friends and family. Where there is an interaction, whether directly with an individual or indirectly with an individual's proxy, there is an opportunity for a threat to manifest itself, resulting in adverse consequences to the individual if conditions permit.

There is a long-standing desire to quantify privacy risks on the part of privacy professionals. Formal assessment of privacy impacts has been a legal requirement in various jurisdictions and contexts around the globe for some time now, not least nowadays by the European Union's General Data Protection Regulation (GDPR) [1]. GDPR Article 25 specifically refers to the likelihood and severity of risks, which, while not exclusive to quantitative approaches, is certainly suggestive of them. Quantitative approaches undoubtedly offer a number of attractive characteristics, including relative ease of summarization and

communication. One of these characteristics is the authoritativeness that derives from quantification, as evinced by other domains such as insurance, which have long exemplified the use of quantitative risk assessment. When done systematically on an appropriate foundation, quantitative risk assessment can be valuable, including as a complement to qualitative approaches. However, while attempts have been made in this vein in the privacy domain, those attempts have been halting and problematic.

These moves have often falsely equated quantification with rigor. As a result, many approaches to quantification of privacy risk have been characterized by a distinct lack of rigor, their use of numerical values notwithstanding. Arguably, a rigorous qualitative analysis is preferable to a non-rigorous quantitative one, and analytical methods exist for the former, e.g., [2]. However, rigorous quantitative approaches to privacy risk analysis are possible as well. This paper describes one such approach, based on Factor Analysis of Information Risk (FAIR), yielding FAIR-Privacy (FAIR-P). While built upon FAIR, FAIR-P incorporates adaptations specific to privacy and is therefore not simply the straightforward application of FAIR to privacy. While the high-level constructs remain the same, the specifics of several of them have been tailored to address privacy. This is particularly the case for consequences, which, for privacy purposes, requires bifurcation to account for both the norm-violative act itself and the potential secondary tangible results of that act.

The remainder of the paper is organized as follows. Section 2 summarizes the history of privacy risk analysis to properly situate the described approach within its larger context. Section 3 provides an overview of FAIR and FAIR-P. Sections 4 and 5 describe proposed quantification, of frequency and severity respectively, within FAIR-P, while Section 6 discusses some of the distinctions between quantitative and qualitative risk assessments. Section 7 reviews three other privacy risk methodologies focused on risks to individuals and compares them against the proposed model. Section 8 offers some concluding thoughts. An appendix in the long version of this paper shows an application of the model to an example case study.

### 2. Short History of Privacy Risk Analysis

Privacy risk analysis is actually a relatively recent way of conceptualizing privacy problems. Owing to the longstanding notion of privacy as a value to be protected, initial efforts to systematically analyze it in the context of

. Applicable material approved for MITRE public release. Distribution unlimited 21-01156-1.

specific laws, projects, and systems leveraged the concept of an impact assessment, exemplified at the time by environmental impact assessments [3]. Privacy impact assessments (PIAs), in turn, leveraged Fair Information Practice Principles (FIPPs), a set of principles which originated in a Code of Fair Information Practices identified in the early 1970s in response to the US Governments collecting and storing massive amounts of data about individuals in government computer systems [4]. Use of PIAs expanded and multiplied into many distinct though quite similar versions, including [5], [6], [7]. GDPR [8] continued this trend. While the impact assessment model, when properly deployed, arguably engendered thoughtful consideration of relevant privacy problems, as constructed it also imposed two fundamental constraints.

The first, and more important, of these constraints has been the reliance on and the embedding of FIPPs in PIA instruments. While it is unsurprising that PIAs seized upon FIPPs as the basis for their structure and analyses, the absence at the time of viable alternatives obscured the fact that a particular choice was being made. As used in PIAs, FIPPs constitute an abbreviated form of risk model. Risk models amount to structured templates of candidate threats and/or vulnerabilities and/or consequences, possibly along with associated factors. (The US National Institute of Standards and Technology (NIST) construes risk models as sets of defined factors [8].) In principle, any impact assessment (or risk management) framework can utilize any relevant risk model. However, when a risk model is embedded in a framework, its role and potential alternatives can effectively become invisible.

The second constraint imposed by the structure adopted by PIAs was the perceived nature of the impact assessment construct itself. Impact assessments (nowadays including social, human rights, and algorithmic assessments) tend to be seen as value-based or value-protective, distinguished from what are perceived as more technocratic risk assessment approaches, such as those used in the context of safety or security. The longstanding view of privacy as being primarily grounded in law rather than science or engineering reinforced this perceived distinction. However, as privacy increasingly manifested as a property of complex socio-technical systems in addition to being a civil liberties issue – as reflected by privacy enhancing technologies (PETs) and Privacy by Design (PbD) – the inadequacies of traditional PIAs and their reliance on FIPPs became evident.

As a result, talk of a “risk-based approach” to privacy arose over the last decade. The proposed instruments of this approach (e.g., [9] and [10]) reflected (usually informally) the perceived inadequacies of FIPPs: the absence of social context, a procedural orientation that served the interests of organizations more than those of individuals, and their lack of any normative grounding. The development of the bases of alternative privacy risk models has enabled augmentation of FIPPs (e.g., Solove’s Taxonomy [11]) and in some cases even replacement (e.g., NIST’s Privacy Risk Assessment Methodology, PRAM [12]). In the latter case, FIPPs, due to the compliance obligations they normally represent, can effectively become requirements checklists rather than a basis for risk analysis. It is also possible to synthesize tailored privacy risk models from multiple existing models, including FIPPs. (See, for example, the

risk model employed in a privacy risk analysis of a proposed architecture for connected vehicles [13].)

While some approaches to privacy risk analysis are quantitative, these have tended toward arbitrary quantification detached from any objective grounding. Typically, the analyst is required to score likelihood and impact on an ordinal scale (often 1 - 10) in the absence of any meaningful defined criteria (e.g., [8]). As a result, such scoring is inevitably arbitrary and inconsistent. Therefore, while such schemes generate numerical values for risk, those values have no intrinsic meaning. There are viable alternatives to this approach to quantification of privacy risk, though. These include the adaptation of FAIR presented here.

### 3. Introduction to FAIR

Factor Analysis of Information Risk is a quantitative risk analysis methodology for calculating information security risk [14]. FAIR breaks down risk into constituent factors which, using probabilistic estimates, can be used to estimate risks arising from information security events. While most readers may be familiar with risk presented in terms of likelihood and impact, FAIR uses a slight variant of that: frequency and magnitude, typically measured on an annualized basis and in dollars, respectively. Hence,  $\text{Risk} = (\text{frequency of loss events}) \times (\text{magnitude of loss})$ .

This paper modifies and expands upon FAIR to introduce a quantitative methodology for assessing privacy risks. Two key features distinguish this privacy version from the original FAIR. First, FAIR looks at organizational risks, whereas FAIR-P (privacy) addresses risks specifically to individuals. Many organizational activities create an external effect on individuals that may not be fully internalized by the organization. Organizational risks can flow from privacy risks to individuals caused by organizational activities, but our analysis principally centers on the externalities imposed on individuals.

This is a key differentiator of the FAIR-P model from other common risk models for credit, operational or insurance risk. Those risk models consider multiple threat actors (e.g., borrowers) against a single at-risk organization (e.g., insurance company). Privacy risk is more akin to safety risk: One or more threat actors against multiple at-risk individuals.

The second distinguishing feature is that where FAIR quantifies magnitude in financial terms, FAIR-P uses a relative scale to rate severity of a privacy harm combined with an additional consideration for the risks of tangible harms flowing from the underlying activity. Since the severity scale has a basis in the non-normativity of the harm, surveys of affected individuals or experimental studies would be the ideal way to scale the harm. Such a survey aligns with the direction of Article 35(9) of the EU General Data Protection Regulation which states “Where appropriate, the controller shall seek the views of data subjects or their representatives on the intended processing, without prejudice to the protection of commercial or public interests or the security of processing operations”. In addition, where tangible adverse consequences on the individual may arise, those are measured in terms appropriate to the consequence such as monetary terms for financial loss and as time for lost liberty. While

some courts may make a habit of translating years of incarceration into a financial figure for compensation for wrongful imprisonment, just because you can assign an arbitrary monetary figure to something does not mean you should. Rather than talking in terms of tradeoffs in dehumanizing dollar figures, tangible impacts should be expressed in appropriate terms (such as deaths, suicides, imprisonments, embarrassments, etc.). Balancing of interests is not always best thought of in purely monetary terms, especially given the externality of the impacts.

In calculating privacy risks, all quantifications in FAIR-P use a probabilistic model based on empirical data and/or appropriate standard distributions, rather than specific numerical value. This is due to uncertainty in assigning specific values to specific factors. If, in fact, any factor is known with certainty, that specific value could be substituted. The values used are also specific to a category of at-risk individuals within whatever system one is examining. We are not looking at specific individuals but rather risks to a group that face similar threats (such as "users of an app" or "bystanders"). Because a probabilistic model does not lend itself to exact calculation, we use, similar to FAIR, a Monte Carlo simulation to envision thousands or millions of trials playing out. Those trials are then plotted on a histogram to show the variance of likely risk values, giving a risk analyst a picture of potential risk.

### 3.1. Harms

While most often associated with physical or mental injury, the term 'harm' can also refer to moral injury or wrongfulness [15]. It is this latter definition with which we concern ourselves. The authors are partial to Daniel Solove's Taxonomy of Privacy harms as it provides a very discrete and granular list of social norms. Readers are free to substitute their own normative framework to identify the harms to be avoided in their analysis. Other models include, but are not limited to, Ryan Calo's Objective and Subjective Harms [16], Alan Westin's four states of privacy [17], Prosser's privacy torts [18], Hartzog's obscurity, trust and autonomy [19], and Nissenbaum's contextual integrity [20]. Work has also been done on systematically deriving synthetic consequences from combined normative models [21].

Solove's taxonomy breaks harms into sixteen discrete harms under four categories, as shown in Table 1. In Solove's original formulation, he prefaced the Collection category with Information similar the first two in the table. We modified that because the underlying harms are not about information or data collected but the act of collecting. Asking an invasive question is a harm regardless of whether the subject provides a response. The act of watching or monitoring someone's activities is the harm regardless of whether data or information is, in fact, collected. Even the appearance of surveillance (a dummy camera), what Ryan Calo calls the "perception of unwanted observation" [16], can result in measurable behavioral changes in subjects.

## 4. Frequency Factors

Within FAIR-P, frequency is a function of four distinct factors: opportunity, motivation, capability, and difficulty.

TABLE 1. SOLOVE TAXONOMY OF PRIVACY HARMS

INFORMATION PROCESSING	INFORMATION DISSEMINATION	COLLECTION	INVASION
Aggregation Identification Insecurity Secondary Use Exclusion	Breach of Confidentiality Increased Accessibility Distortion Disclosure Exposure	Surveillance Interrogation	Decisional Interference Intrusion

We will discuss each of these in turn, then explain how frequency is derived from them.

### 4.1. Opportunity

**Definition:** *Random variable representing the number of opportunities presented to a threat actor in a defined time period (usually one year) to violate the privacy of individuals within an at-risk population.*

When conducting a privacy risk analysis, one must consider what opportunity the threat actor community is presented to impact the at-risk group's privacy. Namely, how frequently will threat actors interact with individuals or interact with proxies for individuals?

For principal threat actors, like a company running a website, determining opportunity becomes a rather straightforward calculation. How many website visitors (the "at-risk" individuals) visit the website in a year? Ideally, one should aim to determine realistic numbers based on empirical data, especially if one has access to website logs. Whether one measures the number of opportunities in terms of unique individuals or unique visits to the website depends upon the nature of the privacy harm being assessed. You can view privacy harms as discrete events or as a continuous harm. Take intrusion, such as spamming, for an example. One view would be where each email a threat actor sends could be an intrusion, and thus the threat actor's opportunities are measured as fast as the threat actor's server will send or as fast as the recipient's client will download the spam. In this view, though, the other factors must be adjusted appropriately (motivation might be not to spam as fast as you can less the recipient block your server). The continuous view would take it, that a threat actor who obtains an email address has "an" opportunity to spam. Whether that means 1 message, 100 message or 1,000,000 messages, by obtaining the email address they now have "an" opportunity.

For indirect threat actors, such as vendors and partners of the website, calculation of opportunity frequency rests on how often they encounter new information about at-risk individuals. For instance, assume a mortgage application website submits the application to different brokers depending on whether the applicant meets certain criteria (location, size of loan). Brokers may only see a fraction of the applications the website operator does.

Opportunity is sometimes described in risk literature as the capacity of a threat actor. A mortgage website has the capacity to threaten the privacy of applicants. However, it does not have the capacity to threaten bus passengers on a bus in Paris. There, the bus driver does.

Given an estimated number of opportunities presented to the threat actors, we could use strictly that number in our calculations. Owing to some uncertainty in the estimate, a distribution around that number would be appropriate. Because opportunity represents discrete events

or actions in a given time frame, a discrete lower bounded distribution is appropriate. Assuming the opportunities are independent, a Poisson distribution makes sense because we know the average number (or at least an estimate) and the outcomes are binary (either the threat actor has an opportunity or they do not) [22].

## 4.2. Motivation

**Definition:** *Random variable representing the likelihood a threat actor will seize an opportunity.*

The next factor to consider is the motive of the threat actors to commit a privacy harm. Motivations vary but we can generalize as done in Table 2.

TABLE 2. THREAT ACTOR CATEGORIES AND COMMON MOTIVES

Persons	Organizations	Governments
Curiosity, spite, money, control, & revenge	Money, competitive advantage, & social change	Law enforcement, espionage, social control, service improvement, & repression

While the above suggests why threat actors might act in a certain way, it does not quantify how motivated the threat actors are. For this, one should consider sub-factors such as threat actors' reward in performing the activity (e.g., making money) and the cost to do so. While financial motivations may be easy to determine, not all threat actors may be rational economic actors, so that should be considered as well. In addition, the risks to threat actors, like being arrested for illegal privacy violations, represents an indirect cost to be considered and any perceived reward discounted against.

Motivation is expressed as a percentage chance of the threat actors taking advantage of opportunities. As with all the calculations, this should be expressed as a probability distribution based on defined scenarios, not a precise numerical value. The probability represents the uncertainty and variation in threat actors' behaviors. A Beta distribution, which works well for distributing percentages, could be used for simulating possible values for motivation. The original FAIR proposed a beta-PERT distribution for its "probability of action" factor, using a confidence value to skew the distribution towards the minimum, maximum or most likely. PERT distributions are commonly used in risk management for handling the uncertainty of estimates [23]. As with the other factors, analysts are free to choose a distribution most suitable to their situation, especially if empirical data is available on threat actor motivations.

## 4.3. Capability

**Definition:** *Random variable representing the strength of the skills and resources available to complete an attempt.*

Capability represents the skills and resources available to threat actors to take advantage of an opportunity. While an employee at a firm may be given access to a database, if they do not have the technical know-how to extract the data or the resources, such as computing power, to use it, they lack the capability to commit the potentially privacy violative act.

Capability is expressed as a percentage chance of the threat actor succeeding when trying to take advantage of an opportunity (absent any impediments – which are introduced below). The probability represents the uncertainty in the threat actor's capability to succeed. Similar to motivation measured as a percentage, a Beta-PERT distribution, continuous over a bounded range between 0 and 1, can be used for simulating possible values for capability.

## 4.4. Difficulty

**Definition:** *Random variable representing the strength of impediments thwarting an attempt*

While one cannot normally affect the capability of a threat actor, one can create impediments that make attacks more difficult for threat actors and raise the threshold at which less capable threat actors are excluded. Consider two threat actors, an 8-year old child and a cryptanalyst employed by a government. Both of them want to read the text messages between the 8-year-old's parents. The skills and resources available to the cryptanalyst might be 1000 times greater than that of the 8-year-old child. We can put the child's most likely capability at 0.099% and the cryptanalyst's at 99.9%. A simple substitution cypher might thwart the 8-year-old but not the cryptanalyst. Because we cannot be certain with precision about the difficulty, we again use a range defining the lower and upper bounds and a most likely estimate for rating the difficulty. As with the other factors, when estimating values reference to objective measures provides some confidence in the numbers being used. In lieu of objective measures, techniques such as calibrated estimates can provide reasonable levels of confidence. (See [24].)

Note that difficulty is distinguished from the deterrent effect of making something difficult. In other words, employing strong encryption may have an effect of making it harder to succeed but also have an effect on motivation, by driving the costs up.

## 4.5. Calculating Frequency

The frequency is derived from attempt frequency and vulnerability. Attempt frequency is a function of opportunity and motivation. In order to derive attempt frequency, a Monte Carlo simulation is performed using both opportunity and motivation. For every simulated opportunity, a value should be picked from the motivation distribution. This value is compared against a random value, over [0,1], representing a threshold with that opportunity to determine if the simulated threat actor was sufficiently motivated to make an attempt. The results represent the distribution of attempts for one simulated trial. Empirical data about attempt frequency could also be used in lieu of calculating attempts based on opportunity and motivation. An analyst may have historical data about threat actor behavior which could be used for this purpose.

Similar to deriving attempt frequency from underlying factors, a Monte Carlo simulation is run against capability and difficulty to determine vulnerability. The distribution represents the uncertainty and variance of threat actors' skills, resources and ultimately successes or failures. An attempt is deemed successful if the random value chosen



over the capability distribution of the threat actor exceeds the random value chosen over the difficulty distribution put in place. Vulnerability is defined as the success rate of these simulated trials. Ideally one would want to calculate a success or failure for each attempt to determine the number of successful attempts (i.e. the frequency of violations for the time period under review) for each simulated trial. In lieu of that, one could calculate an average vulnerability using a Monte Carlo simulation and use that as the success rate in a binomial distribution on attempts to determine frequency for each simulated time period.

## 5. Severity and Tangible Consequences Risks

As indicated in Section 1, risk is an adverse consequence associated with some indicator of severity and a frequency of occurrence. Section 4 covered the frequency of a particular threat manifesting. This section covers the adverse consequences, both in terms of severity of the privacy/moral harm and the secondary risks of more tangible consequences.

### 5.1. Severity of Harm

**Definition:** *Random variable representing the degree of nonnormativity of the privacy violative act relative to other acts constituting the same harm.*

The severity of the harm represents possibly the hardest concept in privacy risk for many to grasp. Most risk frameworks look to tangible harms or damages. We specifically reject an exclusive focus on concrete (physical and mental) harms for two reasons. First, what we aim to prevent are violations of individuals' privacy, not damages. If someone wiretaps your phone, that is a violation of your privacy. It does not matter whether you ever find out about it, or whether the wiretapper uses information learned from the phone call. The act of wiretapping violates social norms, violates expectations of the confidentiality of communications and constitutes an invasion of privacy regardless of whether any subsequent physical or mental injuries befall you. Privacy risk should capture the frequency and severity of this occurring.

Secondly, if the focus of privacy risk is on tangible damages, controls that mitigate these injuries without mitigating the underlying harms would be given credence. If a threat actor places a hidden camera in your house and we are focused on your embarrassment upon finding the camera, the threat actor can make the camera smaller, harder to find, or encrypt the transmission to reduce the chance that you find out and are embarrassed. But those controls have done nothing to reduce the frequency or severity of the underlying violation. In fact, one could argue that increasing the covertness of the surveillance may have made the violation even more severe in society's eyes. By way of example, when it was discovered that the US based retailer Target was inferring pregnancy through in-store purchases, rather than fully disclose such actions to shoppers or simply stop, they obfuscated baby related merchandise in targeted mailers. This reduced the tangible consequences of judgment or ostracism by others in the household but did not alter the underlying "secondary use" of data by Target [25].

The authors propose using the following factors, using a simple mnemonic (ABC), to relatively rate the severity of privacy harm. The factors mirror the Belmont ethical principles which came out of a 1978 commission report on ethical principles for protection of human subjects of research [26].

- **Awareness:** How aware is the individual of the activity? How aware is the individual of why the activity is taking place?
- **Benefit:** How beneficial is the activity to the individual?
- **Consent:** How consensual is the activity? (In measuring consensuality, one must consider any power imbalance between individual and threat actor even in the face of apparent technical consent.)

While the factors are objective, the determination of a numerical range for a specific answer set for the factors would be subjective. However, the hope would be that the same person or organization rating different situations would produce relatively consistent ratings, thus leading to the ability to internally compare privacy risks of different situations within acceptable tolerances and for prioritization.

In considering the probabilistic range of severity, one must also consider the variances in the at-risk populations. Not everyone in the population will have the same awareness, benefit to the same degree or consent at the same level. This will be especially apparent if the at-risk population contains vulnerable sub-populations, such as children. However, severity should not incorporate the risks of adverse consequences. While adverse consequence risk does factor into privacy risk, it should not be doubly incorporated by way of severity.

Why, when the authors criticized the arbitrary quantification in many attempts at privacy risk calculation, do we seem to be introducing it here? First, many of those methods we criticize place both likelihood and impact on an arbitrary scale. With the exception of the severity, all of the other factors presented here have a much more formalized basis and the numbers clearly relate to one another. An opportunity of twice a year is twice as often as an opportunity of once a year, whereas a likelihood of "2" on some ordinal scale just means more likely than a "1" but does not describe how much more likely, failing to provide a meaningful sense of even relative likelihood.

Second, impact ordinal scales reflect the cultural norms of the scales' designers, not necessarily those of the people being affected. The scales are also often tied to data points, like types of data or number of individuals affected, which are at once divorced from context and conflate multiple types of privacy harms into one calculation. Health data is often rated a 10 on an ordinal scale, but a doctor using a patient's pregnancy to determine medical care is worlds away from a retailer inferring it to market to customers, and both of those pale in comparison to a doctor inviting non-medical friends into a delivery room to observe. Severity under FAIR-P is meant to consider context of activities through the three proposed factors of awareness, benefit and consent.

We are not suggesting that estimates of severity should flow solely from the perceptions of the analyst. Using surveys of potentially affected populations would be one

way that severity could be captured to rank and compare the degree to which activities violate social norms of behavior.

## 5.2. Risks of Tangible Consequences

**Definition:** *Product of random variables representing the frequency and magnitude of tangible injuries to the affected population or others.*

**Individuals.** Beyond the amorality of violative acts, as measured by severity, individuals may additionally suffer tangible adverse consequences as a result of the underlying activity. The Future of Privacy Forum’s Harms of Automated Decision-Making [27] is a useable source of potential adverse consequences, it’s particular focus notwithstanding. They can include those affecting the individual’s psyche, such as embarrassment and anxiety, changes in behavior, or more objective external harms, such as financial losses, lost opportunities and lost liberties.

As with any risks, tangible consequences are measured by frequency of occurrences and magnitude within the at-risk population. Not all individuals will suffer embarrassment at having nude photos leaked online and not all individuals will suffer the same degree of embarrassment, for example.

**Society.** Beyond adverse consequences to individuals, privacy harms can introduce broad societal consequences. Privacy is intimately tied to freedom, social mobility, democratic society and individualism. Constant chipping away at individual privacy results in desensitization which can negatively affect other aspects of a free society. While it is beyond the scope of this paper to address all societal harms associated with the erosion of privacy, some aspects include:

**Social Sorting** – categorization of individuals based on shared characteristics with subsequent homogenous treatment where differential treatment is warranted.

**Differential Access** – inequitable treatment of individuals based on distinguishing characteristics where equal treatment is warranted.

**Class Discrimination** – reduced opportunity for social mobility based on social sorting of groups.

## 5.3. Calculating Severity

In a similar fashion to frequency, as discussed in section 4.5, the magnitude factors can be simulated using the Monte Carlo method to model the impacts of random events. For every simulated violation in the trial time period under review, a random value in the distribution of severities should be chosen. The risk for that trial is a summation of simulated severity values for all of the simulated violations – i.e. a product of violation frequency and severity.

While the authors do not specify any particular quantification for severity, we suggest two possible methods:

- Binary (0 or 1): zero represents a non-violation and one indicates the action has exceeded the


threshold to be labeled a violation of privacy. Two issues still need to be addressed. First, the question of labeling something a violation illuminates the tension between a subjective assessment by the data subject and the objective assessment of society. Second, regardless of the line drawn between subjective and objective views of the activity, the heterogeneity of the population will vary the values for each individual.

- Bounded range (0 to 1): similar to the binary approach, with zero and one representing a non-violation and full violation respectively. The continuum in-between represents a degree within a single individual.

## 5.4. Calculating Tangible Consequences Risks

Similar to severity, for each violation in a trial under the Monte Carlo simulation, a tangible consequence risk should be calculated. First, a frequency of those consequences should be calculated. Not every action will result in adverse consequences (e.g. a spam message may be caught by a filter and never require the recipient to delete it). While a threat actor may send out thousands of emails to their list, only a portion will be result in inconvenience. That portion is the frequency of the tangible consequence amongst the potential recipients. The magnitude of the tangible consequences represents the variation of the recipients in the amount of time they spend deleting the messages. Table 3 shows a hypothetical calculation on the inconvenience of spamming. The numbers are meant to be illustrative of the methodology and not representative of any real-world measurements.

TABLE 3. SHOWING 3 TRIALS IN A MONTE CARLO SIMULATION AND HOW TO DERIVE THE RISK OF INCONVENIENCE FROM SPAMMING

Factor	Example	Monte Carlo Simulated value Trial 1	Trial 2	Trial 3	Trial...
Frequency of activity	Spamming	103,029 (meaning threat actors sent out 103,029 spam emails)	74,042	118,379	...
Frequency of Tangible Consequence	Inconvenience	62,458 (meaning 62,458 of those spam emails inconvenienced the recipients)	51,244	93,456	...
Magnitude of Tangible Consequence	How long were they inconvenienced?	For each of the 62,458 times people were inconvenienced, an amount of time is calculated that they were inconvenienced. Chart at right show a Beta-PERT distribution of 0.1 to 5 seconds with most likely of 0.9 seconds.			
Risk of Tangible Consequence	How much inconvenience for the entire population that was at-risk?	41,356 seconds (meaning of the hundred thousand spam messages sent out, the recipients will spend a combined 41,356 seconds dealing with them)	37,004 seconds	82,297 seconds	...

Each type of adverse consequence risks would be expressed separately from each other owing to the heterogeneity in the metrics (financial harm to individuals measured in dollars, lost liberty measured in time, etc.). An analyst could use empirical information to plot a distribution of adverse consequences.

## 6. Qualitative versus Quantitative Analysis

Many qualitative risk models rate risks on a categorical scale (most often low to moderate to high). But these

kinds of models often suffer from the inability to prioritize controls consistent with organizational risk appetite and tolerance. (Exceptions to this include System-Theoretic Process Analysis for Privacy (STPA-Priv) [2].) Many risks will be lumped into the moderate category [13] and some worthwhile controls may not reduce a risk out of a moderate rating. Even a control that reduces a risk from high to moderate does not tell the organization if the associated expense is worthwhile. Additionally, the authors are unaware of any qualitative risk model that explicitly distinguishes between the norm-violative and tangible injuries to individuals.

Quantitative risk analysis seeks to address these deficiencies. If an organization's risk can be measured in dollars and a control that costs \$1 million a year can reduce annual risk by \$10 million, then it is clearly worth the investment. This works well for FAIR applied to information security risks. However, two problems still persist when using a quantitative risk framework for privacy. As previously stated, privacy risks often constitute externalities. There is clearly a financial disincentive to spend money internally to principally benefit those outside the firm. To do so requires conscious organizational motivation and effort. Secondly, not all privacy risks are easily quantified financially. And, if you do quantify embarrassment or lost liberty (such as in years of incarceration), determining risk tolerance for that may be problematic. While some industries, such as transportation, have for years dealt with quantification of injury and death, most companies have not thought very deeply about the effects of their activities in this manner.

Quantification of privacy risk may be methodological overkill in many cases, though, so how does one leverage this privacy risk analysis methodology without massive amounts of calculations and estimations? One way is to use the principal risk factors in this model in a binary fashion, discounting negligible values.

- **Opportunity** – Does what you are doing create an opportunity for a threat actor?
- **Motivation** – Is the threat actor at all motivated?
- **Severity** – Is there general consensus on the rightfulness or wrongfulness of the behavior? (One should escape organizational myopia and look outside for guidance on normative behavior.)
- **Consequences** – Is there the potential for adverse consequences in the at-risk population or society?

One can then use these factors to drive, at a high level, the need for controls to reduce risk. Where one answers affirmatively to all the above necessitates the most attention. Where there is opportunity and motive, this requires the next most direct attention. Finally, where only opportunity exists, the least attention may be paid. In cases where factors exist but organizational inertia and costs prevent action, quantitative calculations can then be utilized to measure benefits of potential expenditures in reducing risk. One could also use the presence of certain factors as thresholds that indicate the need to employ particular types of analysis. Different privacy risk analysis techniques may be more or less appropriate for any particular circumstance. It is helpful to have a toolbox rather than a single tool.

## 7. Comparison

A team from KU Leuven recently proposed a privacy risk methodology (Data-Subject Aware Privacy Risk) [28] which, like FAIR and the proposed method here, breaks privacy risk into frequency and magnitude factors. The frequency factors across all three models are representative of the same concepts: frequency equates to opportunities acted on by a threat actor where the strength of the threat actor can be defeated by countermeasures. Table 4 compares FAIR-P, proposed here, against this methodology as well as NIST's Privacy Risk Assessment Methodology [8] and CNIL's Privacy Risk Methodology [10]. The comparisons are in three areas: what the model seeks to avoid, the measure of likelihood of the avoidable event and the measure of the severity of the avoidable event.

TABLE 4. COMPARISON OF FOUR PRIVACY RISKS ASSESSMENT METHODS

	FAIR-P	Privacy Risk Assessment for Data Subject Aware Threat Modeling [28]	NIST-PRAM [8]	CNIL's Privacy Risks Methodology [10]
To be avoided	Solve Taxonomy (or other normative privacy models) <ul style="list-style-type: none"> <li>• Information Processing</li> <li>• information Dissemination</li> <li>• Collection</li> <li>• Invasion</li> </ul>	Not explicit but discussed "hard privacy" of: <ul style="list-style-type: none"> <li>• Linkability</li> <li>• Identifiability</li> <li>• Detectability</li> <li>• Disclosure of Information</li> </ul>	Problematic Data Actions: <ul style="list-style-type: none"> <li>• Appropriation</li> <li>• Distortion</li> <li>• Induced Disclosure</li> <li>• Insecurity</li> <li>• Re-identification</li> <li>• Stigmatization</li> <li>• Surveillance</li> <li>• Unanticipated Revelation</li> <li>• Unwarranted Restriction</li> </ul>	Fearful events: <ul style="list-style-type: none"> <li>• unavailability of legal processes</li> <li>• change in processing</li> <li>• illegitimate access to personal data</li> <li>• unwanted change in personal data</li> <li>• disappearance of personal data</li> </ul>
Likelihood	Frequency based on: <ul style="list-style-type: none"> <li>• Opportunity</li> <li>• Motivation of threat actor</li> <li>• Threat actor capability</li> <li>• Difficulty of impediments</li> </ul>	Frequency based on: <ul style="list-style-type: none"> <li>• Contact frequency</li> <li>• Probability of action</li> <li>• Threat actor capability</li> <li>• Strength of threat actor or countermeasures bypassed</li> <li>• Retention period</li> </ul>	Arbitrary 10-point ordinal scale	4-point ordinal scale <ol style="list-style-type: none"> <li>1) Negligible</li> <li>2) Limited</li> <li>3) Significant</li> <li>4) Maximum</li> </ol>
Severity	Magnitude based on: <ul style="list-style-type: none"> <li>• Non-normativity of activity</li> <li>• Secondary Consequences Risk</li> </ul>	Magnitude based on <ul style="list-style-type: none"> <li>• # of data subjects</li> <li>• # of records</li> <li>• Data subject type</li> <li>• Data type sensitivity</li> </ul>	Cumulative arbitrary 10-point ordinal scale based on organizational impacts: <ul style="list-style-type: none"> <li>• Noncompliance costs</li> <li>• Direct business costs</li> <li>• Reputational costs</li> <li>• Culture costs</li> </ul>	4-point ordinal scale <ol style="list-style-type: none"> <li>1) Negligible</li> <li>2) Limited</li> <li>3) Significant</li> <li>4) Maximum</li> </ol>
Additional		Uses Monte Carlo simulations with Beta-PERT distributions for all factors.	Prioritization based on two-dimensional plot of likelihood and severity.	Prioritization based on two-dimensional plot of likelihood and severity.

While one minor difference between KU Leuven's model and FAIR-P resides in the former's inclusion of "Retention Period" as a factor influencing the likelihood of an adverse event, it is not much of a distinction since opportunity in FAIR-P subsumes contact frequency and retention in information-centric threats. One difference between the two models plays out in severity, in that the KU Leuven model places the volume of harm into severity whereas FAIR-P incorporates volume in frequency. Consider the threat of "Disclosure of Information" which is the most similar threat under the two models. The hack of iPhone accounts in 2014 resulted in the exposure of nude photos of at least 100 celebrities [29]. FAIR-P would classify this as 100 opportunities to disclose whereas the KU Leuven model would say 100 data subjects impacted (alternatively you could consider it on a per picture basis and suggest 400 opportunities or 100 data subjects with several records per subject, respectively). Regardless, the multiplicative nature of likelihood and severity in both models would equate to a similar risk distribution. The key distinctions between the two models therefore appears to be the data-centricity of the KU Leuven model compared with FAIR-P's broader privacy approach and the

KU Leuven model's lack of explicit consideration of the non-normativity of the activity. The latter seems problematic because without such consideration, disclosure with consent or for the benefit of a data subject would have equivalent risk to disclosure without consent or benefit.

The approaches proposed by NIST and CNIL are similar in that they both use ordinal scales. As previously mentioned, this causes two problems. First, it ignores the breadth and depth of risk variance and uncertainty by consolidating risk into a single quantity or two value coordinates. Secondly, ordinal ranking obscures interstitial variance. An increase in likelihood from a value of 2 to 3 might represent a hundred-fold increase in likelihood while 2 to 3 on the severity scale might represent a mere doubling, but risk as a function of the products of these numerical values would fail to show that.

## 8. Conclusion

The purpose of this paper has been threefold. First, it provides a more formal foundation for assessing quantitative privacy risk as a means of moving beyond less rigorous quantitative approaches. Second, it incorporates in the assessment methodology the external nature of many privacy harms imposed by organizations on individuals and society. Finally, it injects the concept of normative privacy behavior (and not just physical and mental injury) into the equation. Hopefully, this model is useful.

Not addressed in this paper are characteristics and determinations of organizational privacy risk tolerance and appetite. There has been relatively little research in this area. Clearly, more such research is needed to help guide organizations and society in adopting appropriate levels of privacy risk. Additional areas of research include population surveys similar to [30] with comparison scenarios to establish relative degrees of nonnormative behavior.

The appendix to this paper showcases the application of FAIR-P to an example use case of the risks of surveillance by managers of smart locks on home occupants. [31]

## References

- [1] European Union, "Regulation on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (Data Protection Directive)," L119. (Brussels, 2016), 1–88, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679>.
- [2] S. Shapiro, Privacy Risk Analysis Based on System Control Structures: Adapting System-Theoretic Process Analysis for Privacy Engineering, 2016 IEEE Security and Privacy Workshops (SPW), (San Jose, CA, 2016) 17-24.
- [3] R. Clarke, *Comp. L. & Sec. Rev.*, 25, 123-135 (2009), [https://www.researchgate.net/publication/223224533\\_Privacy\\_impact\\_assessment\\_Its\\_origins\\_and\\_development](https://www.researchgate.net/publication/223224533_Privacy_impact_assessment_Its_origins_and_development)
- [4] US Department of Health, Education, and Welfare (HEW), Secretary's Advisory Committee on Automated Personal Data Systems, Records, Computers and the Rights of Citizens. (Department of Justice, Washington, DC, 1973), <https://aspe.hhs.gov/report/records-computers-and-rights-citizens>.
- [5] US Office of Management and Budget (OMB), Managing Information as a Strategic Resource, Circular A-130 (Washington, DC, 2016), <https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/circulars/A130/a130revised.pdf>
- [6] Organization for Economic Development and Cooperation (OECD), The OECD Privacy Framework. (2013), [https://www.oecd.org/sti/economy/oecd\\_privacy\\_framework.pdf](https://www.oecd.org/sti/economy/oecd_privacy_framework.pdf)
- [7] Asia-Pacific Economic Cooperation (APEC), APEC Privacy Framework (2015), APEC#217-CT-01.9. Singapore (2017), [https://www.apec.org/-/media/APEC/Publications/2017/8/APEC-Privacy-Framework-\(2015\)/217\\_ECSG\\_2015-APEC-Privacy-Framework.pdf](https://www.apec.org/-/media/APEC/Publications/2017/8/APEC-Privacy-Framework-(2015)/217_ECSG_2015-APEC-Privacy-Framework.pdf)
- [8] US National Institute of Standards and Technology (NIST), Privacy Risk Assessment Methodology (February 2019), <https://www.nist.gov/document/nist-pram-feb2019zip>.
- [9] Centre for Information Policy Leadership, A Risk-based Approach to Privacy: Improving Effectiveness in Practice. (Washington, DC, 2014), [https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/white\\_paper\\_1-a\\_risk\\_based\\_approach\\_to\\_privacy\\_improving\\_effectiveness\\_in\\_practice.pdf](https://www.informationpolicycentre.com/uploads/5/7/1/0/57104281/white_paper_1-a_risk_based_approach_to_privacy_improving_effectiveness_in_practice.pdf)
- [10] Commission Nationale de l'Informatique et des Libertés (CNIL), Methodology for Privacy Risk Management June, 2012 (in English), <https://www.cnil.fr/sites/default/files/typo/document/CNIL-ManagingPrivacyRisks-Methodology.pdf>
- [11] D. J. Solove, *Understanding Privacy*, 6th edition (Harvard University Press, Cambridge, 2010)
- [12] NIST IR 8062 National Institutes of Standards, An Introduction to Privacy Engineering and Risk Management in Federal Systems, <https://nvlpubs.nist.gov/nistpubs/ir/2017/NIST.IR.8062.pdf>
- [13] US Department of Transportation, Privacy Issues for Consideration by USDOT Based on Review of Preliminary Technical Framework (Final – Rev A), Report Number FHWA-JPO-15-236. (Washington, DC, 2016), <https://www.regulations.gov/document?D=NHTSA-2016-0126-0003>
- [14] J. Jones, J. Freund, *Measuring and Managing Information Risk: A FAIR Approach*, Butterworth-Heinemann (2015)
- [15] Editors of the American Heritage Dictionaries, *The American Heritage Dictionary of the English Language*, 5th edition (Houghton Mifflin Harcourt, 2016)
- [16] R. Calo, *Indiana L. J.*, 86, (2011), [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1641487](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1641487)
- [17] A. Westin, *Wash. & Lee L. Rev.*, 25, (1968), <https://scholarlycommons.law.wlu.edu/wlulr/vol25/iss1/20>
- [18] W. Prosser, *Calif. L. Rev.* 48, 383 (1960)
- [19] W. Hartzog, *Privacy's Blueprint*, (Harvard University Press, 2018)
- [20] H. Nissenbaum, *Wash. L. Rev.* 79, 119-157 (2004)
- [21] S. Shapiro, *Deriving and Using Synthetic Consequences for Privacy and Other Risk Modeling*, (The MITRE Corporation, Bedford, MA, 2020)
- [22] F. A. Haight, *Handbook of the Poisson Distribution*, (John Wiley & Son, 1967)
- [23] D. Vose, *Risk analysis: a quantitative guide*, 3rd edition (John Wiley & Sons, 2008)
- [24] D. Hubbard, *How to Measure Anything: Finding the Value of Intangibles in Business*, 3rd edition (Wiley, 2014)
- [25] K. Hill, *Forbes*, Feb. 16, 2012, <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#3a026e6c6668>
- [26] The Belmont Report Ethical Principles and Guidelines for the Protection of Human Subjects of Research, DHEW (OS) 78-0014, [https://videocast.nih.gov/pdf/ohrp\\_belmont\\_report.pdf](https://videocast.nih.gov/pdf/ohrp_belmont_report.pdf)
- [27] Future of Privacy Forum, *Unfairness by Algorithm, Distilling the Harms of Automated Decision-Making* (2017), <https://fpf.org/2017/12/11/unfairness-by-algorithm-distilling-the-harms-of-automated-decision-making/>
- [28] L. Sion, D. Van Landuyt, K. Wuyts, W. Joosen, 2019 IEEE Security and Privacy Work-shops (SPW), 2019, (KU Leuven, Heverlee, Belgium, 2019) 64-71



- [29] C. Arthur, The Guardian, Sep. 1, 2014. <https://www.theguardian.com/technology/2014/sep/01/naked-celebrityhack-icloud-backup-jennifer-lawrence>
- [30] J. Bhatia, T. D. Breaux, ACM Transactions on Computer-Human Interaction, 10, 25 (2017), <https://www.cs.cmu.edu/~breaux/publications/jbhatia-tochil8.pdf>
- [31] R. Cronk, S. Shapiro, Quantitative Privacy Risk Analysis Appendix <https://enterprivacy.com/risk-tool/Quantitative%20Privacy%20Risk%20Analysis%20Appendix.pdf>

## Appendix A.

### Example Case Study of Surveillance Risks of Smart Locks

For this example, we will discuss how one might approach analyzing the risk of surveillance of a smart lock installed on the exterior of a home (house or apartment) under the FAIR-P model. The at-risk population are the occupants, those occupying the home who will be entering or leaving through the door secured by the smart lock. Assume the “smart” in smart lock refers not just to an electronic lock but something with network capability and the ability to be programmed or managed from that network. The threat actors are those that have administrative capabilities, the “managers” of the lock. This does not imply they act in a management capacity over the occupants, just over the lock. Notice we are not talking about external hackers or someone who needs to breach security to have access to the lock. The concern, expressed by the manufacturer reviewing the privacy risks associated with the smart lock, is that managers, who may also be occupants themselves but could also be landlords or previous occupants, may surveil the occupants, monitoring their comings and goings. This appendix is not intended as a full throttled analysis of the surveillance risks posed by smart locks, but meant to be illustrative of the application of the risk assessment model. We asked Dr. Gilad L. Rosner, Founder of the Internet of Things Privacy Forum, to role play an executive at a hypothetical smart lock manufacturer.

#### A.1. Opportunity

We need three pieces of information to determine threat actor opportunity: the number of threat actors (managers), the number of at-risk individuals (occupants) and the relationship between the two. According to Grand View Research, over 7 million smart locks were sold in 2019. Quite a few of these are used in the hospitality industry or other non-residential settings, though. Dr. Rosner explained that even the residential market is broken up into two very different markets because of the differing technology needs: those selling to single family residences and those selling to managers of multi-family complexes (like apartments). His fictitious company was focused on single family residences and expected to have 500,000 units in place worldwide in the coming year. According to Pew Research the average number of household occupants (in the US) in 2018 was 2.63 (the-number-of-people-in-the-average-u-s-household-is-going-up-for-the-first-time-in-over-160-years/). 2016 numbers from the European Council put the number of household occupants (in the EU) at 2.3. Not knowing the relative market distribution for this company, we average the two numbers to 2.465. However, we’re not concerned about people surveilling themselves, so the number of non-manager occupants at risk would be 1.465. You may be thinking that the managers have an opportunity to surveil the occupants every day, or every time they enter or leave the household. But, by being given administrative rights on the smart lock they are afforded “an” opportunity. Whether they observe the occupant once, once a month, once a day or every five minutes, the observations collectively represent the surveillance. The monitoring is the totality of the activity, not the discrete acts of observance.

$$1 \text{ manager} \times 1.465 \text{ occupants} \times 500,000 \text{ households} = 732,500 \text{ opportunities}$$

- . <https://www.grandviewresearch.com/industry-analysis/smart-lock-market>
- . <https://www.pewresearch.org/fact-tank/2019/10/01/>

The 732,500 opportunities are only an estimate and thus we will use that as the mean in a Poisson distribution to represent uncertainty in our estimation and variability in population throughout the year.

#### A.2. Motivation

A survey was conducted with 192 participants, approximately half in EU and half in the US. Participants were asked a simple question of whether, if they were the lock manager, they would monitor occupants. Participants were asked in two different contexts: as a family member monitoring other family members and apartment manager monitoring tenants. Participants in the US were much more likely to monitor occupants than in the EU, but in both groups, more participants were likely to monitor family members than tenants, though in the EU the number was only marginally higher. A follow-up survey of 91 US based persons asked about familial surveillance was very close (at 61.5%) to the original survey (at 62.4%).

TABLE 5. PARTICIPANTS INTERESTED IN MONITORING FAMILY MEMBERS OR TENANTS

Would you monitor	Participants	Family Members	Tenants
EU	100	38.0%	37.0%
US	92	62.4%	47.4%
<b>Combined</b>	192	49.7%	41.9%

Participants were given an option to comment. Some relevant comments were “I would like to have access, but that does not mean that I would use it all the time, but it can be useful in some situations,” “If I were the owner of a home and used this smart lock, I may find it useful to track when a maid or some other paid service is accessing my home/property,” and “Whether I would want to and whether I actually would are different. I might want to but refrain because of legal fears or ethical qualms. As an apartment manager, it would be mainly for practical reasons (to send maintenance workers in (with the tenant’s permission) when the tenant is away, to avoid disruption).”

In quantifying motivation for analysis, we used the 38.0%, 49.7% (combined) and 62.4% in our Beta-PERT distribution for minimum, most likely and maximum, respectively. See Figure 1. A more sophisticated undertaking might have involved weighting the most likely value to the different geographic markets based on distribution of sales in those markets.

#### A.3. Attempt Frequency

The histogram in figure 2 illustrates the attempt frequency. In a simulation of 2700 possible outcomes, attempts range from a low of 285,315 per year to a high of 452,601. The median was 364,439. Due to limitations in Excel at doing sub-simulations, for each trial, opportunity and motivation were used as inputs in a binomial distribution inversion function as trials and success rate, respectively, to derive the attempt frequency for that trial year.

#### A.4. Vulnerability

We assume, for our baseline risk analysis, that the capability of the managers always exceeds any negligible impediments. This may be an over simplification, as the capabilities of our manager population may include some who lack the technological skills to adequately use the management tool to surveil the population, despite the lack of protective controls. But, for now, the vulnerability plays no part in the calculation.

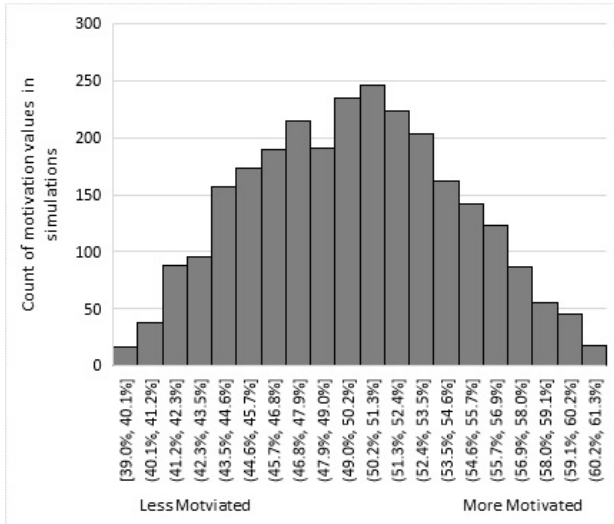


Figure 1. Histogram of simulation counts over possible motivation percentages




Figure 2. Histogram of attempt frequency

### A.5. Frequency

Since vulnerability has no effect on an attempt (every attempt will be successful), Frequency will be the same as attempt frequency. As with attempt frequency, due to limitations of conducting sub-simulations in Excel, one can use attempts and the success rate (i.e. vulnerability) as trials and probability of success in a binomial distribution inversion function to determine successful violations for each simulated trial. This was done for residual risk calculations illustrated in the Risk section below.

### A.6. Severity

First, we need to consider the severity of the harm, surveillance. Survey participants were asked a threshold question of whether they viewed monitoring in the two scenarios as a “privacy violation.” The results are in Table 6. The threshold question purposefully did not intone any awareness, benefit or consent of the occupant. Some participants made potential assumptions about awareness: “I assume that the installation of the smart lock is done with full knowledge of its capabilities to the users.” Others may have seen some countervailing benefit: “I feel like

TABLE 6. WHETHER PEOPLE VIEW SURVEILLANCE IN THE SCENARIO AS A PRIVACY VIOLATION

Is this a privacy violation?	Participants	Family Members	Tenants
EU	100	56.0%	78.0%
US	92	28.3%	90.2%
Combined	192	42.7%	83.8%




Figure 3. Plot of distribution of society’s view as whether monitoring a family member constitutes a privacy violation

I answered yes for the apartment manager/tenant scenario having lived in a dangerous apartment complex, I get why you might want to monitor people coming and going. But for family I feel like it goes too far,” and “It would help if there was a crime and needed to know who was there.” Another participant said “I think anyone out-side of a family monitoring someone usage upon entering their home would be a privacy issue. I can see where parents might want to monitor their children’s exact time of coming and leaving the home. But anything outside that scenario would just be weird”

TABLE 7. TRIALS SHOWING WHICH PERCENT OF ACTIONS FOR EACH TRIAL EXCEED THRESHOLD AND THE RESULTING RISK VALUES INDICATING THE NUMBER OF PRIVACY VIOLATIONS IN A YEAR.

Trial	Frequency	Percentage of actions exceeding threshold	Risk
1	326,410	38.90%	129,967
2	411,510	32.51%	133,798
3	353,869	46.89%	165,917
4	347,492	43.70%	151,837
5	378,232	50.82%	192,222
8	...	...	...

Because we chose to use the binary method measuring severity (i.e. of the activity rising to the threshold of a violation), we used the family members range to in a Beta-PERT distribution (Figure 3) to determine the percentage of actions that would exceed society’s threshold. This is illustrated in Table 7.

## A.7. Risk

We can calculate an annual risk curve at this point (using the Monte Carlo simulations to plot the distributions of potential risk values), but isolated it does not provide much value. If the risks were solely expressible as dollars, the manufacturer could purchase insurance or take other financial account of the risk. But we are talking about a quantification of an externally imposed risk. What we can do is compare this risk to other related values, namely tolerance for risk, residual risk after controls are imposed or similar privacy risks caused by the manufacturer (for use in prioritization of mitigation efforts). The figures below show the results of such comparisons.




Figure 4. Comparison of risk to tolerance. Corporate tolerance can be determined through interviews and surveys with company managers and executives to determine their ranges for acceptable risk. The graph illustrates that risk is significantly higher than risk tolerance. Tolerance was determined by the interview with Dr. Rosner role playing a corporate executive from the company.




Figure 5. Comparison to residual risk after access restrictions to historical logs has been put in place by the manufacturer and only granted in certain circumstances (law enforcement request, etc.) and notification on the locks to occupants about the logs. For illustration we set the access restrictions at 95% effective strength. Note, while better it is still not within tolerance.




Figure 6. Comparison to residual risk after logging was turned off by default and only enabled by the manufacturer for a small set of customers (1%). Now residual risk has been reduced to near tolerance levels.




Figure 7. Comparison of risk profiles for surveillance in the smart lock product to another of the company's products, a smart garage door, illustrating the former has higher baseline risk and thus should take priority in mitigation.