Towards formalizing the GDPR's notion of singling out

Aloni Cohen[a,b,c,1,2] **and Kobbi Nissim**[d,1,2]

[a]Rafik B. Hariri Institute for Computing and Computational Science & Engineering, Boston University, Boston, MA 02215; [b]School of Law, Boston University, Boston, MA 02215; [c]Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139; and [d]Department of Computer Science, Georgetown University, Washington, DC 20007

Edited by Moshe Y. Vardi, Rice University, Houston, TX, and approved March 3, 2020 (received for review August 21, 2019)

There is a significant conceptual gap between legal and mathematical thinking around data privacy. The effect is uncertainty as to which technical offerings meet legal standards. This uncertainty is exacerbated by a litany of successful privacy attacks demonstrating that traditional statistical disclosure limitation techniques often fall short of the privacy envisioned by regulators. We define "predicate singling out," a type of privacy attack intended to capture the concept of singling out appearing in the General Data Protection Regulation (GDPR). An adversary predicate singles out a dataset x using the output of a data-release mechanism *M*(x) if it finds a predicate *p* matching exactly one row in x with probability much better than a statistical baseline. A data-release mechanism that precludes such attacks is "secure against predicate singling out" (*PSO secure*). We argue that PSO security is a mathematical concept with legal consequences. Any data-release mechanism that purports to "render anonymous" personal data under the GDPR must prevent singling out and, hence, must be PSO secure. We analyze the properties of PSO security, showing that it fails to compose. Namely, a combination of more than logarithmically many exact counts, each individually PSO secure, facilitates predicate singling out. Finally, we ask whether differential privacy and *k*-anonymity are PSO secure. Leveraging a connection to statistical generalization, we show that differential privacy implies PSO security. However, and in contrast with current legal guidance, *k*-anonymity does not: There exists a simple predicate singling out attack under mild assumptions on the *k*-anonymizer and the data distribution.

privacy | GDPR | singling out | differential privacy | *k*-anonymity

Data-privacy laws—like the Health Insurance Portability and Accountability Act, the Family Educational Rights and Privacy Act (FERPA), and Title 13 in the United States; and the European Union's (EU's) General Data Protection Regulation (GDPR)—govern the use of sensitive personal information.* These laws delineate the boundaries of appropriate use of personal information and impose steep penalties upon rule breakers. To adhere to these laws, practitioners need to apply suitable controls and statistical disclosure-limitation techniques. Many commonly used techniques, including pseudonymization, k-anonymity, bucketing, rounding, and swapping, offer privacy protections that are seemingly intuitive, but only poorly understood. And while there is a vast literature of best practices, a litany of successful privacy attacks demonstrates that these techniques often fall short of the sort of privacy envisioned by legal standards (e.g., ref. 1).

A more disciplined approach is needed. However, the significant conceptual gap between legal and mathematical thinking around data privacy poses a challenge. Privacy regulations are grounded in legal concepts, such as personally identifiable information (PII), linkage, distinguishability, anonymization, risk, and inference. In contrast, much of the recent progress in data-privacy technology is rooted in formal mathematical privacy models, such as differential privacy (2), that offer a foundational treatment of privacy, with provable privacy guarantees, meaningful semantics, and composability. And while such models are being actively developed and implemented in the academy, industry, and government, there is a lack of discourse between the legal and mathematical conceptions. The effect is uncertainty as to which technical offerings adequately match expectations expressed in legal standards (3).

Bridging between Legal and Technical Concepts of Privacy

We aim to address this uncertainty by translating between the legal and the technical. To do so, we begin with a concept appearing in the law, then model some aspect of it mathematically. With the mathematical formalism in hand, we can better understand the requirements of the law, their implications, and the techniques that might satisfy them.

In particular, we study the concept of "singling out" from the GDPR (4). We examine what it means for a data-anonymization mechanism to ensure "security against singling out" in a data release. Preventing singling out attacks in a dataset is a necessary (but maybe not sufficient) precondition for a dataset to be considered effectively anonymized and thereby free from regulatory restrictions under the GDPR. Ultimately, our goal is to better understand a concept foundational to the GDPR, by enabling a rigorous mathematical examination of whether certain classes of techniques (i.e., k-anonymity and differential privacy) provide security against singling out.

We are not the first to study this issue. The Article 29 Data Protection Working Party, an EU advisory body, provided guidance about the use of various privacy technologies—including k-anonymity and differential privacy—as anonymization techniques (5). Its analysis is centered on asking whether each technology effectively mitigates three risks: "singling out, linkability,

Significance

This article addresses a gap between legal and technical conceptions of data privacy and demonstrates how it can be minimized. The article focuses on "singling out," which is a concept appearing in the GDPR. Our analysis draws on the legislation, regulatory guidance, and mathematical reasoning to derive a technical concept—"predicate singling out"—aimed at capturing a core part of GDPR's intent. Examination of predicate singling out sheds light on the concept of singling out and the question of whether existing technologies protect against such a threat. Conceptually, this work demonstrates the role that principled analysis supported by mathematical argument can and should play in articulating and informing public policy at the interface between law and technology.

Author contributions: A.C. and K.N. designed research, performed research, and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1] A.C. and K.N. contributed equally to this work.

[2] To whom correspondence may be addressed. Email: aloni@bu.edu or kobbi.nissim@georgetown.edu.

First published March 31, 2020.

*Title 13 of the US Code mandates the role of the US Census.

and inference." For instance, their "Opinion on Anonymisation Techniques" concludes that with k-anonymity, singling out is no longer a risk, whereas with differential privacy, it "may not" be a risk (5). Though similar in purpose to our work, its technical analyses are informal and coarse. Revisiting these questions with mathematical rigor, we recommend reconsidering the Working Party's conclusions.

This work is part of a larger effort to bridge between legal and technical conceptions of privacy. An earlier work analyzed the privacy requirements of FERPA and modeled them in a game-based definition, as is common in cryptography (6). The definition was used to argue that the use of differentially private analyses suffices for satisfying a wide range of interpretations of FERPA. An important feature of FERPA that enabled that analysis is that FERPA and its accompanying documents contain a rather detailed description of a privacy attacker and the attacker's goals.

1. Singling Out in the GDPR

We begin with the text of the GDPR (4). It consists of articles detailing the obligations placed on processors of personal data, as well as recitals containing explanatory remarks. Article 1 of the regulation delineates its scope:

> *This regulation lays down rules relating to the protection of natural persons with regard to the processing of personal data and rules relating to the free movement of personal data.*

On the other hand, the GDPR places no restrictions on the processing of nonpersonal data, which includes personal data that have been "rendered anonymous."[†]

The term "personal data" is defined in Article 4 of the GDPR to mean "any information relating to an identified or identifiable natural person; an identifiable natural person is one who can be identified, directly or indirectly." What it means for a person to be "identified, directly or indirectly" is further elaborated in Recital 26:

> *To determine whether a natural person is identifiable, account should be taken of all of the means reasonably likely to be used, such as singling out, either by the controller or by another person to identify the natural person directly or indirectly.*

Thus, singling out is one way to identify a person in data, and only data that do not allow singling out may be excepted from the regulation.

Singling out is the only criterion for identifiability explicitly mentioned in the GDPR, the only occurrence of the term being the passage quoted above. For insight as to the meaning of singling out, we refer to two documents prepared by the Article 29 Data Protection Working Party.[‡] "Opinion on the Concept of Personal Data" (7) elaborates on the meaning of "identifiable, directly or indirectly." A person is identified "within a group of persons [when] he or she is distinguished from all other members of the group." One way of distinguishing a person from a group is by specifying "criteria which allows him to be recognized by narrowing down the group to which he belongs." If the group is narrowed down to an individual, that individual has been singled out.[§] Similarly, the Working Party's "Opinion on Anonymisation Techniques" describes singling out as "isolat[ing] some or all records which identify an individual in [a] dataset." Looking ahead, we will call this "isolating" an individual in the dataset and argue that not every instance of isolation should be considered a singling-out attack.

We highlight three additional insights from ref. 7 that inform our work. First, identification does not require a name or any other traditional identifier. For instance, singling out can be done with a "small or large" collection of seemingly innocuous traits (e.g., "the man wearing a black suit") (7). Indeed, this is what is meant by indirectly identifiable. An example of singling out in practice cited by the Working Party "Opinion on Anonymisation Techniques" (5) showed that four locations sufficed to uniquely identify 95% of people in a pseudonymized dataset of time-stamped locations. This is considered singling out, even without a method of linking the location traces to individuals' names.

Second, identifiable data may come in many forms, including microdata, aggregate statistics, news articles, encrypted data, video footage, and server logs. What's important is not the form of the data—it is whether the data permit an individual to be singled out. We apply this same principle to the manner in which an individual is singled out within a dataset. Most examples focus on specifying a collection of attributes (e.g., four time-stamped locations) that match a single person in the data. A collection of attributes can be viewed as corresponding to a "predicate": a function that assigns to each person in the dataset a value 0 or 1 (interpreted as false or true, respectively). We interpret the regulation as considering data to be personal data if an individual can be isolated within a dataset using any predicate, not only those that correspond to collections of attributes.

Third, whether or not a collection of attributes identifies a person is context-dependent. "A very common family name will not be sufficient to identify someone—i.e., to single someone out—from the whole of a country's population, while it is likely to achieve identification of a pupil in a classroom" (7). Both the prevalence of the name and the size of the group are important in the example and will be important in our formalization.

2. Main Findings

A. Defining Security Against Predicate Singling Out. We formalize and analyze "predicate singling out," a notion which is intended to model the GDPR's notion of singling out. Following the discussion above, we begin with the idea that singling out an individual from a group involves specifying a predicate that uniquely distinguishes the individual, which we call "isolation." Using this terminology, an intuitive interpretation of the GDPR is that rendering data anonymous requires preventing isolation. Trying to make this idea formal, we will see that it requires some refinement.

We restrict our attention to datasets $\mathbf{x} = (x_1, \ldots, x_n)$ of size n, where each row x_i is sampled independently from some underlying probability distribution D over a universe X. The dataset $\mathbf{x}$ is assumed to contain personal data corresponding to individuals, with at most one row per individual. For example, $\mathbf{x}$ might consist of home listings, hospital records, internet browsing history, or any other personal information. A mechanism M takes $\mathbf{x}$ as input and outputs some data release $\mathbf{y} = M(\mathbf{x})$, be it a map of approximate addresses, aggregate statistics about disease, the result of applying machine learning techniques, or pseudonymized internet histories. We call M an "anonymization mechanism" because it purportedly anonymizes the personal data $\mathbf{x}$.

[†] This point is emphasized in Recital 26: "The principles of data protection should therefore not apply to anonymous information, namely information which does not relate to an identified or identifiable natural person or to personal data rendered anonymous in such a manner that the data subject is not or no longer identifiable."

[‡] The Article 29 Data Protection Working Party was established by the EU Data Protection Directive and issued guidance on the meaning of the Directive. The opinions we consider in this article have not been officially updated since GDPR's passing. Moreover, GDPR tacitly endorses the Working Party's guidance on singling out: It borrows the language of Recital 26 almost verbatim from the Data Protection Directive, but adds the phrase "such as singling out."

[§] The notion of "singling out" is not defined in the Opinion on the Concept of Personal Data (7). It is used in ref. 7 four times, each consistent with the above interpretation. Our interpretation coincides with and was initially inspired by that of ref. 8, defining "singling out as occurring when an analyst correctly makes a statement of the form 'There is exactly one user that has these attributes.'"

An adversary A attempts to output a predicate $p: X \to \{0,1\}$ that isolates a row in $\mathbf{x}$, i.e., there exists $i \in [n]$ such that $p(x_i) = 1$ and $p(x_j) = 0$ for all $j \neq i$. We emphasize that it is rows in the original dataset $\mathbf{x}$ on which the predicate acts, not the output $\mathbf{y}$. In part, this is a by-product of our desire to make no assumptions on the form of M's output. While it might make sense to apply a predicate to isolate a row in pseudonymized microdata, it is far from clear what it would mean for a synthetic dataset or for aggregate statistics. Observe that this choice also rules out predicates p that "isolate" rows by referring to their position in $\mathbf{x}$ (i.e., "the seventh row").

M prevents isolation if there doesn't exist an adversary A that isolates a row in $\mathbf{x}$, except with very small probability over the randomness of sampling $\mathbf{x} \leftarrow D^n$, the randomness of the mechanism $\mathbf{y} \leftarrow M(\mathbf{x})$, and the randomness of adversary A$(\mathbf{y})$. Unfortunately, this is impossible to achieve by any mechanism M. To wit, there are trivial adversaries—those that do not look at the outcome $\mathbf{y}$—that isolate a row with probability approximately 37%. The trivial adversaries simply output any predicate p that matches a $1/n$ fraction of the distribution D. Isolation is, hence, not necessarily indicative of a failure to protect against singling out, as a trivial adversary would succeed with $\approx 37\%$ probability (for any n), even if M does not output anything at all.

Example: Consider a dataset of size $n = 365$, including information about random people selected at random from the US population. To isolate a person, a trivial adversary may output $p =$ (born on March 15th). This predicate will isolate a row with probability

$$\binom{365}{1} \cdot \frac{1}{365} \cdot (1 - \frac{1}{365})^{364} \approx 37\%.$$

Furthermore, a trivial adversary need not know the distribution D to isolate with probability $\approx 37\%$, as long as D has sufficient min-entropy (*Fact 3.1*).

A trivial adversary can give us a baseline against which to measure isolation success. But the baseline should not simply be a 37% chance of success. Consider the earlier example of a dataset of 365 random Americans. What if an adversary output predicates like $p =$ (born on March 15th $\wedge$ vegan $\wedge$ speaks Dutch $\wedge$ concert pianist), and managed to isolate 10% of the time? Though 10% is much less than 37%, the predicate is extremely specific and unlikely to isolate a person by chance.

We formalize this intuition by considering the baseline risk of isolation as a function of the weight of the predicate p: the chance that p matches a random row sampled from the distribution D. The baseline for predicates of weight $1/n$ is 37%, but the baseline for an extremely specific predicate may be much lower. The more specific the predicate, the closer the baseline gets to zero. Our primary focus in this paper is on the regime of predicate weights where the baseline is negligible, corresponding to predicates with negligible weight.[¶]

Definition 4.5 (PSO Security, Informal): An adversary predicate singles out a row in $\mathbf{x}$ if it outputs a predicate that isolates a row with probability significantly higher than the baseline risk. A mechanism M is "secure against predicate singling out" (*PSO secure*) if no adversary can use its output $\mathbf{y} = M(\mathbf{x})$ to predicate single out.

B. Analyzing Security Against Predicate Singling Out. Having formulated PSO security, our next goal is to understand the guarantee it offers, what mechanisms satisfy it, and how this concept relates to existing privacy concepts, including differential privacy and k-anonymity.

Two desirable properties of a privacy concept are robustness to "postprocessing" and to "composition." The former requires that if a mechanism M is deemed secure, then anything that can be computed using the outcome of M should also be deemed secure. Hence, the outcome may be reused without creating additional privacy risk. For instance, if a PSO-secure mechanism M outputs microdata, then any statistics that can be computed from those microdata should also be PSO secure. PSO security is robust to postprocessing. The analysis is simple and shown in ref. 9.

Incomposability of PSO Security. We would like that the privacy risk of multiple data releases is not significantly greater than the accumulated risks of the individual releases. In this case, we say that the privacy concept composes. We prove that PSO security does not compose and give two examples of this failure. First, we show that releasing a single aggregate statistic is PSO secure, but superlogarithmically many statistics may allow an adversary to predicate single out. Concretely, a mechanism that outputs a single count is PSO secure. Yet a mechanism that outputs $\omega(\log(n))$ counts may allow an adversary to isolate a row with probability arbitrarily close to one by using a predicate with negligible weight (and negligible baseline). Second, we construct a less natural pair of mechanisms that individually are PSO secure but together allow an adversary to recover a row in the dataset. Consequently, the adversary can predicate single out by isolating this row using a predicate with negligible weight.

Do Differential Privacy and k-Anonymity Guarantee PSO Security? Differential privacy is a requirement of data analyses mechanisms that limits the dependency of a mechanism's output distribution on any single individual's contribution (10). k-anonymity is a requirement from data releases that the (quasi) identifying data of every person in the release should be identical to that of at least $k-1$ other individuals in the release (11, 12) (see *Definitions 6.1* and *7.1* for formal definitions).

Differential privacy is not necessary for PSO security, as an exact count is PSO secure, but is not differentially private. However, differential privacy does provide PSO security. The proof relies on the connection between differential privacy and statistical generalization guarantees (13, 14). We show that predicate singling out implies a form of overfitting to the underlying dataset. If a mechanism is differentially private, it prevents this form of overfitting and, hence, protects against predicate singling out.

On the other hand, we show that k-anonymity does not prevent predicate singling out attacks. Instead, it enables an adversary to predicate single out with probability approximately 37%, even using extremely low-weight predicates for which the baseline risk is negligible. Briefly, the attack begins by observing that typical k-anonymous algorithms "almost" predicate single out. They reveal simple predicates—usually, conjunctions of attributes—that are satisfied by groups of k rows in the dataset. In an effort to make the k-anonymized data as useful as possible, these predicates are as descriptive and specific as possible. To predicate single out a row from the k-anonymous dataset, it roughly suffices to isolate a row from within one of these groups.

C. Implication for the GDPR. Precisely formalizing predicate singling out attacks allows us to examine with mathematical rigor the extent to which specific algorithms and paradigms protect against them. In particular, we show that k-anonymity fails to prevent predicate singling out, but that differential privacy prevents predicate singling out. Our conclusions contrast with those of the Article 29 Working Party: They conclude that k-anonymity eliminates the risk of singling out, while differential privacy "may not" (5). This disagreement may raise doubts about whether our modeling indeed matches the regulators' intent, which we now address.

[¶] For completeness, one can also consider predicates of weight $\omega(\log n/n)$, where the baseline is also negligible. See *Remark 4.2*.

Our goal in interpreting the text of the GDPR and related documents, and in defining predicate singling out, is to provide a precise mathematical formalism to capture some aspect of the concept of personal data (as elucidated in the regulation and in ref. 7) and the associated concept of anonymization. We want to render mathematically falsifiable a claim that a given algorithmic technique anonymizes personal data under the GDPR by providing a necessary condition for such anonymizers.

We argue that predicate singling out succeeds. A number of modeling choices limit the scope of our definition, but limiting the scope poses no issue. Specifically, 1) we only consider randomly sampled datasets; 2) we only consider an attacker who has no additional knowledge of the dataset besides the output of a mechanism; and 3) we do not require that isolation be impossible, instead comparing against a baseline risk of isolation. A technique that purports to anonymize all personal data against all attackers must at least do so against randomly sampled data and against limited attackers. And unless the idea of anonymization mechanisms is completely vacuous, one must compare against a baseline risk.

We don't mean to claim that our modeling is the only one possible. The starting point for the analysis is a description which does not use mathematical formalism, but is, rather, a (somewhat incomplete) description using natural language. Alternative mathematical formalizations of singling out could probably be extracted from the very same text, and we look forward to seeing them emerge.

Policy Implications. Assuming our formalization is in agreement with the GDPR's notion of singling out, the most significant consequence of our analysis is that k-anonymity does not prevent singling out (and likewise ℓ-diversity and t-closeness). Thus, it is insufficient for rendering personal data anonymous and excepting them of GDPR regulation. If PSO security is a necessary condition for GDPR anonymization, then something more is required. On the other hand, differential privacy might provide the necessary level of protection. At least it is not ruled out by our analysis.

More abstractly, we believe that self-composition is an essential property of any reasonable privacy concept. Our finding that PSO security does not self-compose is evidence that self-composition should not be taken for granted, but be a criterion considered when developing privacy concepts.

One may still claim that the assessments made in ref. 5 should be taken as ground truth and that the Article 29 Working Party meant for any interpretation of singling out to be consistent with these assessments. According to this view, the protection provided by k-anonymity implicitly defines the meaning of singling out (partially or in full). Such a position would be hard to justify. To the best of our knowledge, the assessments made by the Article 29 WP were not substantiated by a mathematical analysis. Defining privacy implicitly as the guarantee provided by particular techniques is an approach proven to fail (1).

Is PSO Security a Good Privacy Concept? A predicate singling out attack can be a stepping stone toward a greater harm, even in settings where isolation alone may not. It may enable linking a person's record in the dataset to some external source of information (15), or targeting individuals for differential treatment. As such, it is meaningful as a mode of privacy failure, both in the GDPR context and otherwise.

While we believe that PSO security is relevant for the GDPR as a necessary property of techniques that anonymize personal data, we do not consider it sufficiently protective by itself. First, singling out is only one mode of privacy failure. Many other failure modes have been considered, including reconstruction attacks, membership attacks, inference, and linkage. Second, our definition considers a setting where the underlying data are drawn from some (unknown) underlying distribution, an assumption that is not true in many real-life contexts. In such contexts, PSO security may not prevent singling out under the GDPR. Lastly, the incomposability of PSO security renders it inadequate in isolation and suggests that it should be complemented by other privacy requirements.

3. Notation

A dataset $\mathbf{x} = (x_1, \ldots, x_n)$ consists of n elements taken from the data universe $X = \{0,1\}^d$. We consider datasets where each entry x_i is independently sampled from a fixed probability distribution D over X. We denote by U_d a random variable sampled uniformly at random from $\{0,1\}^d$.

A mechanism M is a Turing machine that takes as input a dataset $\mathbf{x} \in X^n$. Mechanisms may be randomized and interactive.

A predicate is a binary-valued function $p : X \to \{0,1\}$. We define the weight of a predicate p to be $\mathtt{wt}_D(p) \triangleq \Pr_{x\sim D}[p(x) = 1]$. For a dataset $\mathbf{x} \in X^n$, let $p(\mathbf{x}) \triangleq \frac{1}{n}\sum_{i=1}^{n} p(x_i)$. We occasionally use indicator notation $\mathbb{I}()$ to define a predicate; for example, $p(x) = \mathbb{I}(x \in A)$ equals 1 if $x \in A$ and 0 otherwise.

For the purposes of asymptotic analyses, we use the number of rows n in a dataset as the complexity parameter. We omit the dependence of $d = d(n)$ on n.[‖] A function $f(n)$ is negligible, denoted $f(n) = \mathrm{negl}(n)$, if for all $c > 0$, there exists $n_c > 0$ such that $f(n) \le n^{-c}$ for all $n > n_c$. Informally, this means that $f(n)$ approaches 0 faster than any inverse polynomial function for large enough n.

Many of our results apply to all data distributions with sufficient "min-entropy," a measure of randomness useful for cryptography. The min-entropy of a probability distribution D over a universe X is $H_\infty(D) = -\log\left(\max_{y\in X}\Pr_{x\sim D}[x = y]\right)$. The relevant results apply to all distributions with even a moderate amount of min-entropy: $H_\infty(D) > \omega(\log(n)) + \log(1/\alpha)$ for some $\alpha = \mathrm{negl}(n)$ (e.g., $H_\infty(D) > \log^{1+c}(n)$ for $c > 0$). We will call such distributions "min-entropic."

Fact 3.1. *For any min-entropic D and any weight $w \in [0,1]$, there exist predicates p_-, p_+ and a negligible function $\mathrm{negl}(n)$ such that $\mathtt{wt}_D(p_-) \in [w - \mathrm{negl}(n), w]$ and $\mathtt{wt}_D(p_+) \in [w + \mathrm{negl}(n)]$ [using the Leftover Hash Lemma (16)].*

4. Security Against Predicate Singling Out

Consider a data controller who has in its possession a dataset $\mathbf{x} = (x_1, \ldots, x_n)$ consisting of n rows sampled independently from a probability distribution D. We think of the dataset as containing the personal data of n individuals, one per row. The data controller publishes the output of an anonymization mechanism M applied to the dataset $\mathbf{x}$. A predicate singling out adversary A is a nonuniform probabilistic Turing machine with access to the published output $M(\mathbf{x})$ and produces a predicate $p : X \to \{0,1\}$. We abuse notation and write $\mathsf{A}(M(\mathbf{x}))$, regardless of whether M is an interactive or noninteractive mechanism. For now, we assume that all adversaries have complete knowledge of D and are computationally unbounded.[**]

The Article 29 Working Party describes singling out as "isolat[ing] some or all records which identify an individual in [a] dataset" (5). This is done by "narrowing down [to a singleton] the group to which [the individual] belongs" by specifying "criteria which allows him to be recognized" (7). We call this *row isolation*.

Definition 4.1 (Row Isolation): A predicate p isolates a row in a dataset $\mathbf{x}$ if there exists a unique $x \in \mathbf{x}$ such that $p(x) = 1$. That is, if $p(\mathbf{x}) = 1/n$. We denote this event $\mathsf{iso}(p, \mathbf{x})$.

[‖] More formally, we can consider an ensemble of data domains $\mathcal{X} = \{X_n = \{0,1\}^{d(n)}\}_{n\in\mathbb{N}}$ and an ensemble of distributions $\mathcal{D} = \{D_n\}_{n\in\mathbb{N}}$, where each D_n is a distribution over X_n.

[**] As is typical in cryptography, strengthening the adversary to be nonuniform (including possibly having full knowledge of the distribution D) yields stronger security definition. See *Reflections on Modeling*, where we reexamine these choices.

It is tempting to require that an anonymization mechanism M only allows an adversary to isolate a row with negligible probability, but this intuition is problematic. An adversary that does not have access to M—a trivial adversary—can output a predicate p with $\mathsf{wt}_D(p) \approx 1/n$ (by *Fact 3.1*) and thereby isolate a row in $\mathbf{x}$ with probability

$$\binom{n}{1} \cdot \mathsf{wt}_D(p) \cdot (1 - \mathsf{wt}_D(p))^{n-1} \approx (1 - \frac{1}{n})^{n-1} \approx \frac{1}{e} \approx 37\%.$$

In *Bounding the Baseline*, we will see that, in many cases, the trivial adversary need not know the distribution to produce such a predicate.

A. Security Against Predicate Singling Out. We first define a measure of an adversary A's success probability in isolating a row given the output of a mechanism M while being restricted to output a predicate p of at most a given weight w.

Definition 4.2 (Adversarial Success Probability): Let D be a distribution over X. For mechanism M, an adversary A, a dataset size $n \in \mathbb{N}$, and a weight $w \leq 1/n$ let

$$\mathsf{Succ}_w^{\mathsf{A},M}(n, D) \triangleq \Pr_{\substack{\mathbf{x} \leftarrow D^n \\ p \leftarrow \mathsf{A}(M(\mathbf{x}))}} [\mathsf{iso}(p, \mathbf{x}) \wedge \mathsf{wt}_D(p) \leq w].$$

Instead of considering the absolute probability that an adversary isolates a row, we consider the increase in probability relative to a baseline risk: the probability of isolation by a trivial adversary, T.

Definition 4.3 (Trivial Adversary): A predicate singling out adversary T is trivial if the distribution over outputs of T is independent of $M(\mathbf{x})$. That is, $\mathsf{T}(M(\mathbf{x})) = \mathsf{T}(\perp)$.

Definition 4.4 (Baseline): For $n \in \mathbb{N}$ and weight $w \leq 1/n$,

$$\mathsf{base}_D(n, w) \triangleq \max_{\text{Trivial } \mathsf{T}} \mathsf{Succ}_w^{\mathsf{T},\perp}(n, D).$$

The baseline lets us refine the Working Party's conception of singling out as row isolation. We require that no adversary should have significantly higher probability of isolating a row than a trivial adversary, when both output predicates of weight less than w.

Definition 4.5 (Security Against Predicate Singling Out): For $\varepsilon(n) \geq 0$, $\delta(n) \geq 0$, $w_{\mathsf{max}}(n) \leq 1/n$, we say a mechanism M is $(\varepsilon, \delta, w_{\mathsf{max}})$-secure against predicate singling out ($(\varepsilon, \delta, w_{\mathsf{max}})$-PSO secure) if for all A, D, n, and $w \leq w_{\mathsf{max}}$:

$$\mathsf{Succ}_w^{\mathsf{A},M}(n, D) \leq e^{\varepsilon(n)} \cdot \mathsf{base}_D(n, w) + \delta(n).$$

We often omit explicit reference to the parameter n for ε, δ, w_{max}, and to the distribution D when it is clear from context.

The definition is strengthened as ε and δ get smaller. As shown next, $\mathsf{base}(n, w) = \operatorname{negl}(n)$ when $w = \operatorname{negl}(n)$. This is the most important regime of *Definition 4.5*, as such predicates not only isolate a row in the dataset, but likely also isolate an individual in the entire population.

We say a mechanism is PSO secure if for all $w_{\mathsf{max}} = \operatorname{negl}(n)$ there exists $\delta' = \operatorname{negl}(n)$ such that M is $(0, \delta', w(n) w_{\mathsf{max}})$-PSO secure. Observe that for all $\varepsilon = O(\log n)$, $\delta = \operatorname{negl}(n)$, and $w_{\mathsf{max}} \leq 1/n$, $(\varepsilon, \delta, w_{\mathsf{max}})$-PSO security implies PSO security.[††]

B. Bounding the Baseline. In this section, we characterize the baseline over intervals in terms of a simple function $B(n, w)$. For $n \geq 2$ and a predicate p of weight w, the probability over $\mathbf{x} \sim D^n$ that p isolates a row in $\mathbf{x}$ is

$$B(n, w) \triangleq n \cdot w \cdot (1 - w)^{n-1}.$$

$B(n, w)$ is maximized at $w = 1/n$ and strictly decreases moving away from the maximum. $(1 - 1/n)^{n-1}$ approaches e^{-1} as $n \to \infty$ and does so from above (recall that $(1 - 1/n)^n \approx e^{-1}$ even for relatively small values of n).

Claim 4.1. *For all* n, $w \leq 1/n$, *and* D, $\mathsf{base}_D(n, w) \leq B(n, w)$. *For min-entropic* D, $\mathsf{base}_D(n, w) \geq B(n, w) - \operatorname{negl}(n)$.

Proof: Observe that for any randomized trivial adversary T, there exists a deterministic trivial adversary T′ such that $\mathsf{Succ}_w^{\mathsf{T}',\perp}(n) \geq \mathsf{Succ}_w^{\mathsf{T},\perp}(n)$. Therefore, without loss of generality, one can assume that the trivial adversary that achieves the baseline success probability is deterministic.

A deterministic T′ always outputs a predicate p' of weight $\mathsf{wt}(p') = w'$.

$$\Pr_{\mathbf{x} \sim D^n} [\mathsf{iso}(p', \mathbf{x})] = \binom{n}{1} \cdot w' \cdot (1 - w')^{n-1}.$$

Therefore,

$$\mathsf{base}(n, w) = \sup_{\substack{w' \leq w \\ w' = \mathsf{wt}(p')}} n w' (1 - w')^{n-1} \leq B(n, w),$$

where the last inequality follows from the monotonicity of $B(n, w')$ in the range $w' \in [0, w]$. By *Fact 3.1*, there exists p' such that $\mathsf{wt}(p') \geq w - \operatorname{negl}(n)$. Hence, $\mathsf{base}(n, w) \geq B(n, \mathsf{wt}(p')) \geq B(n, w) - \operatorname{negl}(n)$. □

Examples: For a (possibly randomized) function f, consider the mechanism M_f that on input $\mathbf{x} = (x_1, \ldots, x_n)$ outputs $M_f(\mathbf{x}) = (f(x_1), \ldots, f(x_n))$. Whether M_f is PSO secure depends on the choice of f:

- **Identity function:** If $f(x) = x$, then $M_f(\mathbf{x}) = (x_1, \ldots, x_n)$ and every distinct $x_i \in \mathbf{x}$ can be isolated by a predicate $p_i : x \mapsto \mathbb{I}(x = x_i)$. If $\Pr_{x \sim D}[x = x_i] = \operatorname{negl}(n)$ (e.g., when D is uniform over $\{0,1\}^{\omega(\log(n))}$ or D is min-entropic), then $\mathsf{wt}_D(p_i) = \operatorname{negl}(n)$. Hence, M_f is not PSO secure.
- **Pseudonymization:** If f is one-to-one and public, it offers no more protection than the identity function. For unique $x_i \in \mathbf{x}$ and $y_i = f(x_i)$, the predicate $p_i : x \mapsto \mathbb{I}(f(x) = y_i)$ isolates x_i. If D is min-entropic, $\mathsf{wt}_D(p_i) = \operatorname{negl}(n)$. Observe that M_f is not PSO secure, even if f^{-1} is not efficiently computable. Furthermore, f being many-to-one does not guarantee PSO security. For instance, suppose the data are uniform over $\{0,1\}^n$ and $f : \{0,1\}^n \to \{0,1\}^{n/2}$ outputs the last $n/2$ bits of an input x. M_f is not PSO secure. In fact, it is possible to single out every row $x \in \mathbf{x}$ using the same predicates p_i as above. Together, these observations challenge the use of some forms of pseudonymization.
- **Random function:** If $f(x)$ is a secret, random function, then $M_f(\mathbf{x})$ carries no information about $\mathbf{x}$ and provides no benefit to the adversary over a trivial adversary. For every $\mathbf{x}$, a trivial adversary T can perfectly simulate any adversary $\mathsf{A}(M_f(\mathbf{x}))$ by executing A on a random input. Hence, M_f is PSO secure.

C. Reflections on Modeling. In many ways, *Definition 4.5* demands the sort of high level of protection typical in the foundations of cryptography. It requires a mechanism to provide security for all distributions D and against nonuniform, computationally unbounded adversaries.[‡‡] The main weakness in the required

[††] For any $w'(n) = \operatorname{negl}(n)$, let $\delta'(n) = e^{\epsilon(n)}\mathsf{base}(n, w'(n)) + \delta(n) = \operatorname{negl}(n)$.

[‡‡] It is reasonable to limit the adversary in *Definition 4.5* to polynomial time. If we restricted our attention to distributions with moderate min-entropy, the results in this paper would remain qualitatively the same: Our trivial adversaries and lower bounds are all based on efficient and uniform algorithms; our upper bounds are against unbounded adversaries. Relatedly, restricting to min-entropy distributions would allow us to switch the order of quantifiers of D and T in the definition of the baseline without affecting our qualitative results.

protection is that it considers only data that are independent and identically distributed (i.i.d.), whereas real-life data cannot generally be modeled as i.i.d.

Any mechanism that purports to be a universal anonymizer of data under the GDPR—by transforming personal data into nonpersonal data—must prevent singling out. Our definition is intended to capture a necessary condition for a mechanism to be considered as rendering data sufficiently anonymized under the GDPR. Any mechanism that prevents singling out in all cases must prevent it in the special case that the data are i.i.d. from a distribution D and for $w_{\mathsf{max}} = \mathrm{negl}(n)$. We view a failure to provide security against predicate singling out (*Definition 4.5*) as strong evidence that a mechanism does not prevent singling out as conceived of by the GDPR.

On the other hand, satisfying *Definition 4.5* is not sufficient for arguing that a mechanism renders data anonymized under the GDPR. Singling out is only one of the many "means reasonably likely to be used" (4) to identify a person in a data release. Furthermore, the definition considers only i.i.d. data; it may not even imply that a mechanism prevents singling out in other relevant circumstances.

It is important that our definitions are parameterized by the weight w. An unrestricted trivial adversary can isolate a row with probability of about $1/e \approx 37\%$ by outputting a predicate of weight about $1/n$. If 37% was used as the general baseline (without consideration of the weight of the predicate), then the definition would permit an attacker to learn very specific information about individuals in a dataset, as long as it does so with probability less than 37%. For instance, such a definition would permit a mechanism that published a row from the dataset with a one in three chance.

Remark 4.2. *The baseline is also negligible when the trivial adversary is required to output predicates with weight at least $\omega(\log n/n)$. It is not clear to the authors how beneficial finding a predicate in this regime may be to an attacker. This high-weight regime is analyzed analogously in ref. 9.*

5. Properties of PSO Security

Two desirable properties of privacy concepts are 1) immunity to postprocessing (i.e., further processing of the outcome of a mechanism, without access to the data, should not increase privacy risks), and 2) closure under composition (i.e., a combination of two or more mechanisms which satisfy the requirements of the privacy concept is a mechanism that also satisfies the requirements, potentially with worse parameters). It is easy to see that PSO security withstands postprocessing. However, it does not withstand composition, and we give two demonstrations for this fact.

First, we consider mechanisms which count the number of dataset rows satisfying a property and show that every such mechanism is PSO secure. However, there exists a collection of $\omega(\log(n))$ counts which allows an adversary to isolate a row with probability arbitrarily close to one using a predicate with negligible weight. Second, we construct a (less natural) pair of mechanisms that are individually PSO secure, but together allow the recovery of a row in the dataset. This construction borrows ideas from ref. 17.

Not being closed under composition is a significant weakness of the notion of PSO security. Our constructions rely on very simple mechanisms that we expect would be considered as preventing singling-out attacks (as a legal matter), even under other formulations of the concept. As such, it may well be that nonclosure under composition is an inherent property of the concept of singling out.

As a policy matter, we believe that closure under composition (as well as immunity to postprocessing) should be considered prerequisites for any privacy concept deemed sufficient to protect individuals' sensitive data. Pragmatically, the fact that PSO security is not closed under composition suggests that this concept is best used for disqualifying privacy technology (i.e., if it is not PSO secure). This concept should not be used alone to certify or approve the use of any technology.

A. A PSO-Secure Counting Mechanism.

For any predicate $q : X \to \{0,1\}$, we define the corresponding Counting Mechanism:

Algorithm 1: Counting Mechanism $M_{\#q}$:

```
input: $\mathbf{x} = (x_1, \ldots, x_n)$
return $|\{1 \le i \le n : q(x_i) = 1\}|$
```

For example, consider the least-significant bit predicate lsb, that takes as input a string $x \in \{0,1\}^d$ and outputs the first bit $x[1]$. The corresponding Counting Mechanism $M_{\#\mathsf{lsb}}$ returns the sum of the first column of $\mathbf{x}$.

$M_{\#q}$ is PSO secure for any predicate q. This is a corollary of the following proposition.

Proposition 5.1. *For all* A, $n > 0$, $w \le 1/n$, *and* $M : X^n \to Y$, $\mathsf{Succ}_w^{\mathsf{A},M}(n) \le |Y| \cdot \mathsf{base}(n, w)$.

Proof: We define a trivial adversary T such that for all A, $\mathsf{Succ}_w^{\mathsf{T},\perp}(n) \ge \frac{1}{|Y|} \cdot \mathsf{Succ}_w^{A,M}(n)$. The proposition follows by the definition of $\mathsf{base}(n, w)$. T samples a random $y \in_R Y$ and returns $p \leftarrow \mathsf{A}(y)$. For all datasets $\mathbf{x}$, there exists $y^* = y^*(\mathbf{x}) \in Y$ such that

$$\Pr_{p \leftarrow \mathsf{A}(y^*)}[\mathsf{iso}(p, \mathbf{x}) \wedge \mathsf{wt}(p) \le w]$$
$$\ge \Pr_{p \leftarrow \mathsf{A}(M(\mathbf{x}))}[\mathsf{iso}(p, \mathbf{x}) \wedge \mathsf{wt}(p) \le w].$$

For all $\mathbf{x}$, $\Pr_{y \in_R Y}[y = y^*] = \frac{1}{|Y|}$. Therefore,

$$\mathsf{Succ}_w^{\mathsf{T},\perp}(n) = \Pr_{\substack{\mathbf{x} \leftarrow D^n \\ y \in_R Y \\ p \leftarrow \mathsf{A}(y)}}[\mathsf{iso}(p, \mathbf{x}) \wedge \mathsf{wt}(p) \le w] \ge \frac{\mathsf{Succ}_w^{A,M}(n)}{|Y|}.$$

□

As exact counts are not differentially private, the PSO security of $M_{\#q}$ demonstrates that differential privacy (*Definition 6.1*) is not necessary for PSO security. The security of a single exact count easily extends to $O(1)$-many counts (even adaptively chosen), as the size of the codomain grows polynomially.

B. Failure to Compose.

Our next theorem states that a fixed set of $\omega(\log(n))$ counts suffices to predicate single out with probability close to e^{-1}. For a collection of predicates $Q = (q_0, \ldots, q_m)$, let $M_{\#Q}(\mathbf{x}) \triangleq (M_{\#q_0}(\mathbf{x}), \ldots, M_{\#q_m}(\mathbf{x}))$. Let $D = U_d$ be the uniform distribution over $X = \{0,1\}^d$ for some $d = \omega(\log(n))$.

Theorem 5.2. *For any $m \le d$, there exists a Q and* A *such that*

$$\mathsf{Succ}_{2^{-m}}^{\mathsf{A},M_{\#Q}}(n) \ge B(n, 1/n) - \mathrm{negl}(n).$$

Choosing $m = \omega(\log(n))$ yields $2^{-m} = \mathrm{negl}(n)$.

Proof: Treating $x \in \{0,1\}^d$ as a binary number in $[0, 2^d - 1]$, let $q_0(x) = \mathbb{I}(x < 2^d/n)$. Observe that $\mathsf{wt}(q_0) = 1/n - \mathrm{negl}(n)$.

For $i \in \{1, \ldots, m\}$, define the predicate $q_i(x) \triangleq (q_0(x) \wedge x[i])$, and let $y_i = M_{\#q_i}(\mathbf{x})$. Observe that if it happens that q_0 isolates row j^* of $\mathbf{x}$, then $y_i = x_{j^*}[i]$. Consider the deterministic adversary A that on input $M_{\#Q}(\mathbf{x}) = (y_0, \ldots, y_m)$ outputs the predicate

$$p(x) = q_0(x) \wedge \left(\bigwedge_{i=1}^{m} (x[i] = y_i)\right).$$

Observe that $\mathsf{iso}(q_0, \mathbf{x}) \implies \mathsf{iso}(p, \mathbf{x})$ and that by construction $\mathsf{wt}(p) \le 2^{-m}$. Thus,

$$\begin{aligned}\mathsf{Succ}_{2^{-m}}^{\mathsf{A},M_{\#Q}}(n) &= \Pr_{\substack{\mathbf{x}\leftarrow U_d^n\\ p\leftarrow \mathsf{A}(M_{\#Q}(\mathbf{x}))}}[\mathsf{iso}(p,\mathbf{x})]\\ &\geq \Pr_{\substack{\mathbf{x}\leftarrow U_d^n\\ p\leftarrow \mathsf{A}(M_{\#Q}(\mathbf{x}))}}[\mathsf{iso}(q_0,\mathbf{x})]\\ &\geq B(n,1/n)-\mathrm{negl}(n).\end{aligned}$$

□

Remark 5.3. *When the attack succeeds, all of the predicates* q_i *match 0 or 1 rows in* $\mathbf{x}$*. It may seem that an easy way to counter the attack is by masking low counts, a common measure taken, e.g., in contingency tables. However, it is easy to modify the attack to only use predicates* Q *matching* $\Theta(n)$ *rows using one extra query. This means that restricting the mechanism to suppress low counts cannot prevent this type of attack. Let* q^* *be a predicate with* $\mathsf{wt}_{U_m}(q^*)=1/2$ *(e.g., parity of the bits), and let* $q_i^*=q_i\vee q^*$*. The attack succeeds whenever* $q^*(\mathbf{x})=q_0^*(\mathbf{x})+1$*. If* $q^*(x)$ *and* $q_0(x)$ *are independent, then this occurs with probability at least* $\frac{1}{2}\cdot B(n,1/n)-\mathrm{negl}(n)$*. As before, the probability can be amplified to* $1-\mathrm{negl}(n)$*.*

While a single count is PSO secure for any data distribution, the above attack against $\omega(\log(n))$ counts applies only to the uniform distribution U_d. We can extend the attack to general min-entropic distributions D at the cost of randomizing the attacked mechanism M (i.e., the set of predicates Q). Furthermore, for min-entropic D, this probability can be amplified to $1-\mathrm{negl}(n)$ by repetition (9).

C. Failure to Compose Twice. Borrowing ideas from refs. 17 and 18, we construct two mechanisms M_{ext} and M_{enc} which are individually PSO secure (for arbitrary distributions), but which together allow an adversary to single out with high probability when the data are uniformly distributed over the universe $X=\{0,1\}^d$. With more work, this composition attack can be extended to more general universes and to min-entropic distributions.

Theorem 5.4. M_{ext} *and* M_{enc} *described below are PSO secure. For* $m=\omega(\log(n))$ *and* $m\leq\min(d,(n-1)/4)$, $X=\{0,1\}^d$*, and* $D=U_d$ *the uniform distribution over* X*, there exists an adversary* A *such that*

$$\mathsf{Succ}_{2^{-m}}^{\mathsf{A},M_{\mathsf{ExtEnc}}}(n)\geq 1-\mathrm{negl}(n),$$

where $M_{\mathsf{ExtEnc}}=(M_{\mathsf{ext}},M_{\mathsf{enc}})$*.*

We divide the input dataset into two parts. We treat $\mathbf{x}_{\mathsf{ext}}=(x_1,\ldots,x_{n-1})$ as a "source of randomness" and x_n as a "messsage." $M_{\mathsf{ext}}(\mathbf{x})$ outputs an encryption secret key s based on the rows in $\mathbf{x}_{\mathsf{ext}}$, using the von Neumann extractor as described in *Algorithm 2*.

Algorithm 2 $M_{\mathsf{ext}}(\mathbf{x})$*:*

```
s ← $\emptyset$, the empty string;
for $i \leftarrow 1$ to $n-1$ by 2 do
    if $\mathsf{lsb}(x_i)=0 \ \wedge\ \mathsf{lsb}(x_{i+1})=1$ then $\mathsf{s}\leftarrow s\|0$;
    if $\mathsf{lsb}(x_i)=1 \ \wedge\ \mathsf{lsb}(x_{i+1})=0$ then $\mathsf{s}\leftarrow s\|1$;
end
if $|\mathsf{s}|\geq m$ then return $\mathsf{s}[1:m]$, the first $m$ bits of s;
else return $\perp$;
```

$M_{\mathsf{enc}}(\mathbf{x})$ runs $\mathsf{s}\leftarrow M_{\mathsf{ext}}$. If $\mathsf{s}\neq\perp$, it outputs $\mathsf{c}=\mathsf{s}\oplus x_n[1:m]$ (using s as a one-time pad to encrypt the first m bits of x_n); otherwise, it outputs $\perp$. Alone, neither s nor c allows the adversary to single out, but using both, an adversary can recover $x_n[1:m]$ and thereby predicate single out the last row.

Proof Outline: We must prove that M_{ext} and M_{enc} are PSO secure and that M_{ExtEnc} is not. Note that the security of M_{ext} and M_{enc} do not follow merely from the fact that their outputs are nearly uniform.[§§]

M_{ext} is $(\ln(2),0,1/n)$-PSO secure. Consider the mechanism $M_{\mathsf{ext}}^{\sigma}(\mathbf{x})$ that samples a random permutation $\sigma:[n]\to[n]$ and outputs $M_{\mathsf{ext}}(\sigma(\mathbf{x}))$. For any $\mathbf{x}$, $M_{\mathsf{ext}}^{\sigma}(\mathbf{x})$ is uniform conditioned on not outputting $\perp$. Its security is equivalent to a mechanism outputting a single bit of information, which itself is $(\ln(2),0,1/n)$-PSO secure by *Proposition 5.1* (and therefore also PSO-secure by the observation at the end of *Security Against PSO*). To complete the proof, one can show that $\mathsf{Succ}_w^{\mathsf{A},M_{\mathsf{ext}}}(n)=\mathsf{Succ}_w^{\mathsf{A},M_{\mathsf{ext}}^{\sigma}}(n)$. □

M_{enc} is PSO-secure. We separately consider the two possible values of $p(x_n)$, where p is the predicate returned by A. Let $\mathsf{Succ}_w^{\mathsf{A},M_{\mathsf{enc}}}(n)=\gamma_0+\gamma_1$, where

$$\gamma_b\triangleq\Pr\left[\mathsf{iso}(p,\mathbf{x})\wedge\mathsf{wt}_D(p)\leq w\wedge p(x_n)=b\right].$$

The output of $M_{\mathsf{enc}}(\mathbf{x})$ is deterministic and information-theoretically independent of x_n. Thus, for any $w=\mathrm{negl}(n)$,

$$\gamma_1\leq\Pr_{\mathbf{x},\mathsf{A}}[p(x_n)=1\mid\mathsf{wt}_D(p)\leq w]\leq w=\mathrm{negl}(n).$$

If A singles out and $p(x_n)=0$, then it is effectively singling out against the subdataset $\mathbf{x}_{-n}=(x_1,\ldots,x_{n-1})$. That is,

$$\gamma_0=\Pr_{\mathbf{x},\mathsf{A}}[\mathsf{iso}(p,\mathbf{x}_{-n})\wedge\mathsf{wt}_D(p)\leq w\wedge p(x_n)=0].$$

We construct B that tries to single out against mechanism M_{ext} using A. On input s, B samples $x_n'\sim D$ and runs $p\leftarrow\mathsf{A}(\mathsf{s}\oplus x_n'[1:m])$.

$$\begin{aligned}\mathsf{Succ}_w^{\mathsf{B},M_{\mathsf{ext}}}(n)&\geq\Pr\left[\mathsf{iso}(p,\mathbf{x}_{-n})\wedge\mathsf{wt}_D(p)\leq w\wedge p(x_n')=0\right]\\ &\quad\cdot\Pr[p(x_n)=0\mid\mathsf{wt}_D(p)\leq w]\\ &\geq\gamma_0\cdot(1-\mathrm{negl}(n)).\end{aligned}$$

By the PSO security of M_{ext}, γ_0 is negligible. □

Insecurity of M_{ExtEnc} for $D=U_d$. The output of $M_{\mathsf{ExtEnc}}(\mathbf{x})$ is a pair (s,c). If $(\mathsf{s},\mathsf{c})=(\perp,\perp)$, A aborts. By a Chernoff Bound, for $m\leq(n-1)/4$, $\Pr_{\mathbf{x}}[\mathsf{s}=\perp]\leq e^{-(n-1)/16}=\mathrm{negl}(n)$. If $(\mathsf{s},\mathsf{c})\neq(\perp,\perp)$, A recovers $x_n=\mathsf{c}\oplus\mathsf{s}$ and outputs the predicate

$$p(x)=(x[1:m]=x_n[1:m]).$$

By the choice of $m=\omega(\log(n))$, $\mathsf{wt}_{U_d}(p)=2^{-m}<\mathrm{negl}(n)$. $\Pr[\mathsf{iso}(p,\mathbf{x})\mid\mathsf{s}\neq\perp]=1-\Pr[\exists j\neq n:x_j[1:m]=x_n[1:m]]=1-n\cdot 2^{-m}>1-\mathrm{negl}(n)$. The bound on $\mathsf{Succ}_{2^{-m}}^{\mathsf{A},M_{\mathsf{ExtEnc}}}$ follows, completing the proof of the claim and the theorem. □

D. Singling Out and Failure to Compose. The failure to compose demonstrated in *Theorem 5.2* capitalizes on the use of multiple counting queries. Such queries underlie a large variety of statistical analyses and machine learning algorithms. We expect that other attempts to formalize security against singling out would also allow counting queries. If so, our negative composition results may extend to other formalizations.

The failure to compose demonstrated in *Theorem 5.4* is more contrived. We expect that other attempts to formalize security against singling out would allow mechanisms like M_{ext}—ones where for every input $\mathbf{x}$ randomly permuting the input and applying the mechanism results in the uniform distribution over outputs (as in the proof of *Theorem 5.4*). It is less clear to us whether other possible formalizations of security against singling out would allow a mechanism like M_{enc}. If it is to compose, it likely must not.

[§§] For example, the mechanism that outputs x_1 may be uniform, but it trivially allows singling out. Security would follow if the output was nearly uniform conditioned on $\mathbf{x}$, but M_{ext} does not satisfy this extra property.

6. Differential Privacy Provides PSO Security

In this section, we demonstrate that differential privacy implies PSO security. Because exact counts are not differentially private but are PSO secure (*Proposition 5.1*), we have already shown that PSO security does not imply differential privacy.

Recall that we model $\mathbf{x} \in X^n$ as containing personal data of n distinct individuals. For $\mathbf{x}, \mathbf{x}' \in X^n$, we write $\mathbf{x} \sim \mathbf{x}'$ if $\mathbf{x}$ and $\mathbf{x}'$ differ on the data of exactly one individual x_i.

***Definition 6.1* [*Differential Privacy* (10, 19)]** A randomized mechanism $M : X^n \to T$ is (ε, δ)-differentially private if for all $\mathbf{x}, \mathbf{x}' \in X^n, \mathbf{x} \sim \mathbf{x}'$ and for all events $S \subseteq T$,

$$\Pr[M(\mathbf{x}) \in S] \leq e^{\varepsilon} \Pr[M(\mathbf{x}') \in S] + \delta,$$

where the probability is taken over the randomness of the mechanism M.

Our analysis relating PSO security to differential privacy is through a connection of both concepts to statistical generalization. For differential privacy, this connection was established in refs. 13 and 14. We use a variant of the latter from ref. 20:

Lemma 6.1 (Generalization Lemma). *For all distributions D and for all (ε, δ)-differentially private algorithms $\mathsf{A} : \mathbf{x} \mapsto p$ operating on a dataset $\mathbf{x}$ and outputting a predicate $p : X \to \{0, 1\}$*

$$\mathop{\mathbb{E}}_{\mathbf{x} \sim D^n}\left[\mathop{\mathbb{E}}_{p \leftarrow \mathsf{A}(\mathbf{x})}[p(\mathbf{x})]\right] \leq e^{\varepsilon} \cdot \mathop{\mathbb{E}}_{\mathbf{x} \sim D^n}\left[\mathop{\mathbb{E}}_{p \leftarrow \mathsf{A}(\mathbf{x})}[\mathsf{wt}_D(p)]\right] + \delta.$$

Theorem 6.2. *For all $\varepsilon = O(1)$, $\delta = \text{negl}(n)$, and $w_{\mathsf{max}} \leq 1/n$, if M is (ε, δ)-differentially private, then M is $(\varepsilon', \delta', w_{\mathsf{max}})$-PSO secure for*

$$\varepsilon' = \varepsilon + (n-1)\ln\left(\frac{1}{1 - w_{\mathsf{max}}}\right) \leq \varepsilon + 1 \text{ and } \delta' = n\delta + \text{negl}(n).$$

In particular, for $w_{\mathsf{max}} = o(1/n)$, $\varepsilon' = \varepsilon + o(1)$.

Proof: For simplicity of exposition, we present the proof for min-entropic distributions D. The proof for general distributions follows from a similar argument.

Given $p \leftarrow \mathsf{A}(M(\mathbf{x}))$, $w \leq w_{\mathsf{max}}$, and D, define the predicate p^*:

$$p^*(x) \equiv \begin{cases} p(x) & \text{if } \mathsf{wt}_D(p) \leq w \\ 0 & \text{if } \mathsf{wt}_D(p) > w \end{cases}.$$

Observe that $\mathsf{wt}_D(p^*) \leq w$. The predicate p^* can be computed from p, D, and w without further access to $\mathbf{x}$. Because differential privacy is closed under postprocessing, if M is (ε, δ)-differentially private, then the computation that produces p^* is (ε, δ)-differentially private as well. Recall $\mathsf{iso}(p, \mathbf{x}) \iff p(\mathbf{x}) = 1/n$.

$$\begin{aligned}
\mathsf{Succ}_w^{\mathsf{A},M}(n) &= \Pr_{\mathbf{x},p}[p(\mathbf{x}) = 1/n \wedge \mathsf{wt}_D(p) \leq w] \\
&\leq \Pr_{\mathbf{x},p}[p(\mathbf{x}) \geq 1/n \wedge \mathsf{wt}_D(p) \leq w] \\
&= \Pr_{\mathbf{x},p}[p^*(\mathbf{x}) \geq 1/n] \\
&\leq n \cdot \mathop{\mathbb{E}}_{\mathbf{x},p}[p^*(\mathbf{x})] \\
&\leq n \cdot (e^{\varepsilon} w + \delta) && \text{(Lemma 6.1)} \\
&= e^{\varepsilon} \frac{B(n, w)}{(1-w)^{n-1}} + n\delta \\
&\leq e^{\varepsilon} \frac{\mathsf{base}(n, w)}{(1-w)^{n-1}} + n\delta + \text{negl}(n) && \text{(Claim 4.1)} \\
&\leq e^{\varepsilon'} \mathsf{base}(n, w) + \delta'.
\end{aligned}$$

□

7. Does *k*-Anonymity Provide PSO Security?

k-anonymity (11, 12) is a strategy intended to help a data holder "release a version of its private data with scientific guarantees that the individuals who are the subjects of the data cannot be re-identified while the data remain practically useful" (12). It is achieved by making each individual in a data release indistinguishable from at least $k - 1$ individuals along certain attributes. Typically, a k-anonymized dataset is produced by subjecting it to a sequence of generalization and suppression operations.

In this section, we analyze the extent to which k-anonymity provides PSO security. We show that a k-anonymized dataset typically provides an attacker information sufficient to PSO with constant probability. This result challenges the determination of the Article 29 Working Party.[¶¶]

A. Preliminaries. Let $(A_1, \ldots, A_m)$ be attribute domains. A dataset $\mathbf{x} = (x_1, \ldots, x_n)$ is a collection of rows $x_i = (a_{i,1}, \ldots, a_{i,m})$, where $a_{i,j} \in A_j$. For subsets $\widehat{a}_{i,j} \subseteq A_j$, we view $y_i = (\widehat{a}_{i,1}, \ldots, \widehat{a}_{i,m})$ as a set in the natural way, writing $x_i \in y_i$ if $\forall j \in [m]$, $a_{i,j} \in \widehat{a}_{i,j}$. We say that a dataset $\mathbf{y} = (y_1, \ldots, y_n)$ is derived from $\mathbf{x}$ by generalization and suppression if $\forall i \in [n]$, $x_i \in y_i$. For example, if (A_1, A_2, A_3) correspond to "5-Digit ZIP Code," "Gender," and "Year of Birth," then it may be that $x_i = (91015, F, 1972)$ and $y_i = (91010\text{–}91019, \star, 1970\text{–}1975)$, where $\star$ denotes a suppressed value.

k-anonymity aims to capture a sort of anonymity in a (very small) crowd: A data release $\mathbf{y}$ is k-anonymous if any individual row in the release is indistinguishable from $k - 1$ other individual rows. Let $\mathsf{count}(\mathbf{y}, y) \triangleq |\{i \in [n] : y_i = y\}|$ be the number of rows in $\mathbf{y}$ which agree with y.[***]

Definition 7.1 (k-Anonymity [Rephrased from Ref. 12]): For $k \geq 2$, a dataset $\mathbf{y}$ is k-anonymous if $\mathsf{count}(\mathbf{y}, y_i) \geq k$ for all $i \in [n]$. An algorithm is called a k-anonymizer if on an input dataset $\mathbf{x}$ its output is a k-anonymous $\mathbf{y}$ which is derived from $\mathbf{x}$ by generalization and suppression.

For k-anonymous dataset $\mathbf{y} = (y_1, \ldots, y_n)$, let $\phi_1(x) = \mathbb{I}(x \in y_1)$ be the predicate that returns 1 if y_1 could have been derived from x by generalization and suppression. Let $\mathbf{x}_{\phi_1} = \{x \in \mathbf{x} : \phi_1(x) = 1\}$. We assume for simplicity that k-anonymizers always output $\mathbf{y}$ such that $|\mathbf{x}_{\phi_1}| = k$, but the theorem generalizes to other settings.

B. *k*-Anonymity Enables Predicate Singling Out. Before presenting the main theorem of this section, we provide an example of a very simple k-anonymizer that fails to provide security against predicate singling out. Let $D = U_n$ be the uniform distribution over $U = \{0, 1\}^n$.

Consider the k-anonymizer that processes groups of k rows in index order and suppresses all bit locations where any of the k rows disagree. Namely, for each group $g = 1, \ldots, n/k$ of k rows $(x_{gk+1}, \ldots, x_{gk+k})$), it outputs k copies of the string $y_g \in \{0, 1, \star\}^n$, where $y_g[j] = b \in \{0, 1\}$ if $x_{gk+1}[j] = \cdots = x_{gk+k}[j] = b$ (i.e., all of the k rows in the group have b as their jth bit) and $y_g[j] = \star$ otherwise.

The predicate $\phi_1(x)$ evaluates to 1 if $y_1[j] \in \{x[j], \star\}$ for all $j \in [n]$ and evaluates to 0 otherwise. Namely, $\phi_1(x)$ checks whether x agrees with y_1 (and, hence, with all of $x_1, \ldots, x_k)$) on all nonsuppressed bits.

In expectation, $n/2^{k-1}$ positions of y_1 are not suppressed. For large enough n, with high probability over the choice of $\mathbf{x}$, at least

[¶¶] Our results hold equally for ℓ-diversity (21) and t-closeness (22), which the Article 29 Working Party also concludes prevent singling out.

[***] Often count is paramaterized by a subset Q of the attribute domains called a quasi-identifier. This parameterization does not affect our analysis, and we omit it for simplicity.

$\frac{n}{2 \cdot 2^{k-1}}$ positions in y_1 are not suppressed. In this case, $\mathsf{wt}_D(\phi_1) \leq 2^{-\frac{n}{2^k}}$ which is $\mathrm{negl}(n)$ for any constant k.

We now show how ϕ_1 can be used adversarially. In expectation, $n(1 - 2^{-(k-1)}) \geq n/2$ positions of y_1 are suppressed. For large enough n, with high probability over the choice of $\mathbf{x}$ at least $n/4$ of the positions in y_1 are suppressed. Denote these positions $i_1, \ldots, i_{n/4}$. Define the predicate $p_k(x)$ that evaluates to 1 if the binary number resulting from concatenating $x[i_1], x[i_2], \ldots, x[i_{n/4}]$ is greater than $2^{n/4}/k$ and 0 otherwise. Note that $\mathsf{wt}_D(p_k) \approx 1/k$ and, hence, p_k isolates within group 1 with probability $\approx 1/e \approx 0.37$, as was the case with the trivial adversary described after *Definition 4.1*.

An attacker observing ϕ_1 can now define a predicate $p(x) = \phi_1(x) \wedge p_k(x)$. By the analysis above, $\mathsf{wt}(p)$ is negligible (as it is bounded by $\mathsf{wt}(\phi_1)$) and $p(x)$ isolates a row in $\mathbf{x}$ with probability ≈ 0.37. Hence, the k-anonymizer of this example fails to protect against singling out.

Theorem 7.1 captures the intuition from this example and generalizes it, demonstrating that k-anonymity does not typically protect against predicate singling out.

Theorem 7.1. *For any $k \geq 2$, there exists an algorithm* A *such that for all min-entropic D, all k-anonymizers* Anon, *and all $w \leq 1/n$:*

$$\mathsf{Succ}_w^{\mathsf{A},\mathsf{Anon}}(n) \geq \eta \cdot B(k, 1/k) - \mathrm{negl}(n) \approx \frac{\eta}{e},$$
$$\text{where} \quad \eta \triangleq \Pr_{\substack{\mathbf{x} \leftarrow D^n \\ \mathbf{y} \leftarrow \mathsf{Anon}(\mathbf{x})}} [\mathsf{wt}_D(\phi_1) \leq w].$$

To predicate single out, the A must output a predicate that both isolates $\mathbf{x}$ and has low weight. The theorem states that these two requirements essentially decompose: η is the probability that the predicate k-anonymizer induces a low-weight predicate ϕ_1, and $B(k, 1/k)$ is the probability that a trivial adversary predicate singles out the subdataset $\mathbf{x}_{\phi_1}$ of size k. Algorithms for k-anonymity generally try to preserve as much information in the dataset as possible. Thus, we expect such algorithms to typically yield low-weight predicates ϕ_1 and correspondingly high values of η.

Proof Outline: A will construct some predicate q and output the conjunction $p \triangleq \phi_1 \wedge q$. Noting that $\mathsf{wt}_D(p) \leq \mathsf{wt}_D(\phi_1)$, and that $\mathsf{iso}(q, \mathbf{x}_{\phi_1}) \implies \mathsf{iso}(p, \mathbf{x})$,

$$\begin{aligned} \mathsf{Succ}_w^{\mathsf{A},\mathsf{Anon}}(n) &\geq \Pr\left[\mathsf{iso}(q, \mathbf{x}_{\phi_1}) \wedge \mathsf{wt}_D(\phi_1) \leq w\right] \\ &= \eta \cdot \Pr\left[\mathsf{iso}(q, \mathbf{x}_{\phi_1}) \,\middle|\, \mathsf{wt}_D(\phi_1) \leq w\right]. \end{aligned} \qquad [1]$$

It remains only to show that for min-entropic distributions D, $\Pr\left[\mathsf{iso}(q, \mathbf{x}_{\phi_1}) \,\middle|\, \mathsf{wt}_D(\phi_1) \leq w\right] \geq B(k, 1/k) - \mathrm{negl}(n)$. This claim is reminiscent of *Fact 3.1*, but with an additional challenge. The rows in $\mathbf{x}_{\phi_1}$ are not distributed according to D; instead, they are a function of Anon and the whole dataset $\mathbf{x}$. They are not independently distributed, and even their marginal distributions may be different from D. Nevertheless, for the purposes of the baseline, the rows in $\mathbf{x}_{\phi_1}$ have enough conditional min-entropy to behave like random rows. □

Data Availability

This paper does not include original data.

ACKNOWLEDGMENTS We thank Uri Stemmer, Adam Sealfon, and anonymous reviewers for their helpful technical suggestions. This material is based upon work supported by the US Census Bureau under Cooperative Agreement CB16ADR0160001 and by NSF Grant CNS-1413920. A.C. was supported by the 2018 Facebook Fellowship and Massachusetts Institute of Technology's (MIT's) RSA Professorship and Fintech Initiative. This work was done while A.C. was at MIT, in part while visiting Georgetown University, and while both authors visited the Simons Institute at the University of California, Berkeley. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of funders.

1. P. Ohm, Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Rev.* **57**, 1701–1777 (2010).
2. C. Dwork, F. McSherry, K. Nissim, A. Smith, Calibrating noise to sensitivity in private data analysis. *J. Priv. Confid.* **7**, 17–51 (2017).
3. K. Nissim, A. Wood, Is privacy *privacy*? *Philos. Trans. R. Soc. A.* **376**, 20170358 (2018).
4. European Parliamentand the Council of the European Union, Regulation (EU) 2016/679, 206. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32016R0679. Accessed 1 April 2019.
5. Article 29 Data Protection Working Party, Opinion 05/2014 on anonymisation techniques. https://iapp.org/media/pdf/resource_center/wp136_concept-of-personal-data_06-2007.pdf. Accessed 1 April 2019.
6. K. Nissim *et al.*, Bridging the gap between computer science and legal approaches to privacy. *Harv. J. Law Technol.* **31**, 687–780 (2018).
7. Article 29 Data Protection Working Party, Opinion 04/2007 on the concept of personal data. https://iapp.org/media/pdf/resource_center/wp216_Anonymisation-Techniques_04-2014.pdf (2007). Accessed 1 April 2019.
8. P. Francis *et al.*, Extended Diffix-Aspen. arXiv:1806.02075 (6 June 2018).
9. A. Cohen, K. Nissim, Towards formalizing the GDPR's notion of singling out. arXiv:1904.06009 (12 April 2019).
10. C. Dwork, F. McSherry, K. Nissim, A. Smith, "Calibrating noise to sensitivity in private data analysis" in *Third Theory of Cryptography Conference*, S. Halevi, T. Rabin, Eds. (Lecture Notes in Computer Science, Springer, Berlin, Germany, 2006), vol. 3876, pp. 265–284.
11. P. Samarati, L. Sweeney, "Generalizing data to provide anonymity when disclosing information" in *Proceedings of the 17th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, A. O. Mendelzon, J. Paredaens, Eds. (ACM Press, New York, NY, 1998), p. 188.
12. L. Sweeney, k-anonymity: A model for protecting privacy. *Internat. J. Uncertainty Fuzziness Knowledge-Based Syst.* **10**, 557–570 (2002).
13. C. Dwork *et al.*, "Preserving statistical validity in adaptive data analysis" in *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015*, R. A. Servedio, R. Rubinfeld, Eds. (ACM, New York, NY, 2015), pp. 117–126.
14. R. Bassily *et al.*, "Algorithmic stability for adaptive data analysis" in *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016*, D. Wichs, Y. Mansour, Eds. (ACM, New York, NY, 2016), pp. 1046–1059.
15. A. Narayanan, V. Shmatikov, "Robust de-anonymization of large sparse datasets" in *IEEE Symposium on Security and Privacy, 2008, SP 2008* (IEEE, Piscataway, NJ, 2008), pp. 111–125.
16. Y. Dodis, R. Ostrovsky, L. Reyzin, A. Smith, Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *SIAM J. Comput.* **38**, 97–139 (2008).
17. K. Nissim, A. D. Smith, T. Steinke, U. Stemmer, J. Ullman, "The limits of post-selection generalization" in *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS'18)*, S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, Eds. (Curran Associates Inc., Red Hook, NY, 2008), vol. 31, pp. 6402–6411.
18. C. Dwork, M. Naor, On the difficulties of disclosure prevention in statistical databases or the case for differential privacy. *J. Priv. Confid.* **2**, 93–107 (2010).
19. C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, M. Naor, "Our data, ourselves: Privacy via distributed noise generation" in *Annual International Conference on the Theory and Applications of Cryptographic Techniques* (Springer, Cham, Switzerland, 2006), pp. 486–503.
20. K. Nissim, U. Stemmer, On the generalization properties of differential privacy. arXiv:1504.05800 (22 April 2015).
21. A. Machanavajjhala, D. Kifer, J. Gehrke, M. Venkitasubramaniam, "*L*-diversity: Privacy beyond *k*-anonymity" in *22nd International Conference on Data Engineering (ICDE'06)* (IEEE Computer Society, Washington, DC, 2007), p. 1.
22. N. Li, T. Li, S. Venkatasubramanian, "t-closeness: Privacy beyond k-anonymity and l-diversity" in *Proceedings of the 23rd International Conference on Data Engineering, ICDE 2007*, R. Chirkova, A. Dogac, M. T. Özsu, T. K. Sellis, Eds. (IEEE Computer Society, Washington, DC, 2007), pp. 106–115.