ARTICLE

WHAT A HYBRID LEGAL-TECHNICAL ANALYSIS TEACHES US ABOUT PRIVACY REGULATION: THE CASE OF SINGLING OUT

MICAH ALTMAN, ALONI COHEN, KOBBI NISSIM, ALEXANDRA WOOD*

ABSTRACT

This article advocates a hybrid legal-technical approach to the evaluation of technical measures designed to render information anonymous in order to bring it outside the scope of data protection regulation. The article demonstrates how such an approach can be used for instantiating a key anonymization concept appearing in the EU General Data Protection Regulation (GDPR) – singling out. The analysis identifies and addresses a tension between a common, compelling theory of singling out and a mathematical analysis of this theory, and it demonstrates how to make determinations regarding the sufficiency of specific technologies for satisfying regulatory requirements for anonymization.

*Authors are listed in alphabetical order. All authors contributed to conceptualization and writing through commentary, review, editing, and revision. Micah Altman is Director of Research at the Center for Research in Equitable and Open Scholarship at the Massachusetts Institute of Technology, e-mail: escience@mit.edu. Aloni Cohen is a Postdoctoral Associate at the Hariri Institute for Computing at Boston University and the Boston University School of Law, e-mail: aloni@bu.edu. Kobbi Nissim is a McDevitt Chair in Computer Science at Georgetown University and affiliated with Georgetown University Law Center, e-mail: kobbi.nissim@georgetown.edu. Alexandra Wood is a Fellow at the Berkman Klein Center for Internet & Society at Harvard University, e-mail: awood@cyber.law.harvard.edu. The authors wish to thank Michel José Reymond for his many insightful comments on the manuscript. The authors also wish to thank their collaborators with the Privacy Tools Project, the members of the Bridging Privacy Definitions Working Group, and participants at the 2020 Privacy Law Scholars Conference (PLSC 2020) for their thoughtful feedback. Work of K.N. and A.W. was partially supported by the U.S. Census Bureau under cooperative agreement no. CB16ADR0160001 and by a gift to the McCourt School of Public Policy and Georgetown University. Work of A.C. was partially supported by the 2018 Facebook Fellowship and MIT's RSA Professorship and Fintech Initiative. This material is based upon work supported by the National Science Foundation under Grant Nos. CNS-1413920, CNS-1414119, and CNS-1915763, and by DARPA under Agreement No. HR00112020021. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of their funders.

Doubts about the feasibility of effective anonymization and de-identification have gained prominence in recent years in response to high-profile privacy breaches enabled by scientific advances in privacy research, improved analytical capabilities, the wider availability of personal data, and the unprecedented richness of available data sources. At the same time, privacy regulations recognize the possibility, at least in principle, of data anonymization that is sufficiently protective so as to free the resulting (anonymized) data from regulation. As a result, practitioners developing privacy enhancing technologies face substantial uncertainty as to the legal standing of these technologies. More fundamentally, it is not clear how to make a determination of compliance even when the tool is fully described and available for examination.

This gap is symptomatic of a more general problem: legal and technical approaches to data protection have developed in parallel, and their conceptual underpinnings are growing increasingly divergent. When lawmakers rely on purely legal concepts to engage areas that are affected by rapid scientific and technological change, the resulting laws, when applied in practice, frequently create substantial uncertainty for implementation, provide contradictory recommendations in important cases, disagree with current scientific technical understanding, and fail to scale to the rapid pace of technological development. This article argues that new hybrid concepts, created through technical and legal co-design, can inform practices that are practically complete, coherent, and scalable.

As a case study, the article focuses on a key privacy-related concept appearing in Recital 26 of the GDPR called singling out. We identify a compelling theory of singling out that is implicit in the most persuasive guidance available, and demonstrate that the theory is ultimately incomplete. We then use that theory as the basis for a new and mathematically-rigorous privacy concept called predicate singling-out. Predicate singling-out sheds light on the notion of singling out in the GDPR, itself inextricably linked to anonymization. We argue that any data protection tool that purports to anonymize arbitrary personal data under the GDPR must prevent predicate singling-out. This enables, for the first time, a legally- and mathematically-grounded analysis of the standing of supposed anonymization technologies like k-anonymity and differential privacy.[1]

Conceptually, our analysis demonstrates that a nuanced understanding of baseline risk is unavoidable for a theory of singling out based on current regulatory guidance. Practically, it identifies previously unrecognized failures of anonymization. In particular, it demonstrates that some k-anonymous mechanisms may allow singling out, challenging the prevailing regulatory guidance.

The article concludes with a discussion of specific recommendations for both policymakers and scholars regarding how to conduct a hybrid legal-technical

[1] The analysis in this article is backed by a technical-mathematical analysis previously published by two of the authors. *See* Aloni Cohen & Kobbi Nissim, *Towards Formalizing the GDPR's Notion of Singling Out*, 117 PROC. OF THE NAT'L ACAD. OF SCI. 8344 (2020).

analysis. Rather than formalizing or mathematizing the law, the article provides approaches for wielding formal tools in the service of practical regulation.

I. INTRODUCTION

There is continuing debate regarding the effectiveness of anonymization and de-identification for protecting individual privacy. Doubts about the feasibility of effective anonymization have gained prominence in recent years in response to high-profile privacy breaches enabled by scientific advances in privacy research, improved analytical capabilities, the wider availability of personal data,

and the unprecedented richness of the data available. At the same time, privacy regulations recognize the possibility, at least in principle, of data anonymization that is sufficiently protective so as to free the resulting (anonymized) data from regulation. The dissonance raises broad questions about the extent to which public and regulatory conceptions of privacy are logically coherent and practically feasible.

In light of this tension, practitioners developing privacy enhancing technologies face substantial uncertainty regarding the legal standing of these technologies. As a result, developers are investing in designing technical tools for regulatory compliance that may later be shown to fail to provide strong privacy protection in practice and thereby arguably fail to satisfy regulatory requirements.[2] More fundamentally, it is not clear how to make a determination of compliance even when the tool is fully described and available for examination.

This gap between privacy technologies and regulation is symptomatic of a more general problem. Legal and technical approaches to data protection have developed in parallel, and their conceptual underpinnings are growing increasingly divergent. When lawmakers rely on purely legal concepts to engage areas that are affected by rapid scientific and technological change, the resulting laws, when applied in practice, frequently create substantial uncertainty for implementation, provide contradictory recommendations in important cases, disagree with current scientific technical understanding, and fail to scale to the rapid pace of technological development.

The gap between legal and formal technical concepts has negative consequences not only for those developing new privacy technologies but also for regulators evaluating their legal compliance and for society at large. Because legal concepts of privacy are not aligned with the state-of-the-art technical understandings in the field, regulatory standards are forced to continually respond and adapt as new vulnerabilities are discovered in the wild.[3] The uncertainty created by a continually evolving interpretation of a legal requirement can slow the adoption of technologies that offer practical reduction in harm, which is a loss for individuals and society as a whole. While ambiguity of concepts may not be a problem when a general legal concept is refined by application to new domains, ambiguity becomes a problem when it masks inherent contradictions

[2] *See* Aloni Cohen & Kobbi Nissim, *Linear Program Reconstruction in Practice*, 10 J. OF PRIVACY AND CONFIDENTIALITY 1, 2 (2020) (demonstrating that a tool "advertised as an off-the-shelf, GDPR-compliant privacy solution" that had been certified by CNIL as "deliver[ing] GDPR-level anonymity" is vulnerable to reconstruction attacks).

[3] As one example, the US Office of Management and Budget's guidance on protecting personally identifiable information has evolved over time to address new understandings of how de-identified data can be vulnerable to attacks. The latest update to the guidance, paralleling the GDPR's approach to personal information, advises government agencies that they must consider that non-personally identifiable information may become personally identifiable information in the future. *See* OFFICE OF MGMT. AND BUDGET, EXEC. OFFICE OF THE PRESIDENT, MEMORANDUM M-17-12, PREPARING FOR AND RESPONDING TO A BREACH OF PERSONALLY IDENTIFIABLE INFORMATION, 5-8 (2017).

or necessitates the reinterpretation of a practice that was previously accepted widely. Further, the reactive mode of adapting the regulation to address new vulnerabilities (analogous to the "penetrate-and-patch" approach to patching software as new bugs are discovered) can leave acute privacy risks unaddressed for extended periods.

This article illustrates a hybrid legal-technical approach to the development and evaluation of the analysis of technical measures to render personal information anonymous or de-identified, in order to bring them outside the scope of regulations. As a case study, the article focuses on a key privacy-related concept appearing in Recital 26 of the General Data Protection Regulation (GDPR) called *singling out*.

This article identifies a compelling theory of singling out as isolation, which is implicit in the most persuasive guidance available, and demonstrates that the theory is ultimately incomplete. This theory is then used as the starting point for a new and mathematically-rigorous privacy concept called *predicate singling-out*. Predicate singling-out sheds light on the notion of singling out in the GDPR, itself inextricably linked to anonymization. The article argues that any data protection tool that purports to anonymize arbitrary personal data under the GDPR must prevent predicate singling-out. This enables, for the first time, a legally- and mathematically-grounded analysis of the standing of supposed anonymization technologies like k-anonymity and differential privacy.[4]

The analysis in this article has implications for policymakers. Conceptually, it demonstrates that a nuanced understanding of baseline risk is unavoidable for a theory of singling out based on current regulatory guidance. Practically, it can help detect previously unrecognized failures of anonymization. In particular, it demonstrates that some k-anonymous mechanisms may allow singling out, challenging the prevailing regulatory guidance.

A hybrid legal-technical approach is necessary for regulatory frameworks to adapt to the scale of technical change, and for formal analysis to be of practical value. New hybrid concepts, created through technical and legal co-design, can help inform practices that are effectively complete, coherent, and scalable. However, developing hybrid legal-technical concepts is substantially more challenging than developing a concept that functions solely in one area. The article concludes with a discussion of specific recommendations for both policymakers and scholars regarding how to conduct a hybrid legal-technical analysis.

The goal of the article is not to advocate formalizing or mathematizing the law, but rather to provide approaches for wielding formal tools in the service of regulation and practice. We conjecture that development of new privacy concepts that are both practically justiciable and technically coherent requires not just review by both legal and technical experts, but *co-design*. When new cross-domain ("hybrid") concepts are needed, it is vital that experts in all implicated fields be enlisted to assist early in the design process. This is particularly

[4] The analysis in this article is backed by a technical-mathematical analysis previously published by two of the authors of this article. *See* Cohen & Nissim, *supra* note 1, at 8346.

important when dealing with regulatory areas that involve informational, environmental, and systems risk. Without thoughtful regulation in place, individually innocuous actions can accumulate into unexpected, catastrophic, and irreversible damage.[5]

A. Organization

Part II of this article defines ANONYMIZATION and singling out, based on the text of the GDPR and the opinions of the Article 29 Data Protection Working Party.[6] It identifies and rebuts a simplistic theory of singling out. Part III presents the new concept of predicate singling-out and its relationship to GDPR singling-out. It explains and justifies the many modeling choices made along the way and examines some of the strengths and weaknesses of the concept. Part IV uses predicate singling-out to evaluate two popular approaches to data anonymization—*k*-anonymity and differential privacy—identifying a disagreement with the Working Party opinion. Part V discusses the implications of predicate singling-out for the GDPR, especially in light of the disagreement, and suggests model regulatory language. Part VI reflects on the hybrid legal-technical analytical approach more broadly, making recommendations for policymakers and scholars.

Parts II–IV include gray call-out boxes which should be read as part of the surrounding text. Those labeled "Example" are illustrations of ideas that are important to understanding the article and are meant to serve as frequent references for the reader. Those labeled "Technical Details" provide additional mathematical insight for the interested reader and may be skipped at first reading.

B. A hybrid legal-technical approach

This article introduces "hybrid legal-technical concepts," argues that hybrid concepts are necessary for effective regulation of new information technologies in complex domains, and demonstrates how hybrid concepts can be developed using multi-disciplinary modes of theorizing.

Although both the term and idea proposed here are (to the best of our knowledge) novel, this work builds on and intersects with a number of existing areas of study. The most senior of these is the economic analysis of law, which

[5] *See* Aaron Fluitt, Aloni Cohen, Micah Altman, Kobbi Nissim, Salome Viljoen & Alexandra Wood, *Data Protection's Composition Problem*, 3 EUR. DATA PROTECTION L. REV. 285, 285-86 (2019).

[6] *See* Regulation (EU) 2016/679, of the European Parliament and of the Council of 27 April 2016 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 (General Data Protection Regulation) (2016) O.J. L119/1, Recital 26 [hereinafter GDPR]; ARTICLE 29 DATA PROTECTION WORKING PARTY, OPINION 04/2007 ON THE CONCEPT OF PERSONAL DATA 21 (2007) [hereinafter *Working Party Opinion on Personal Data*]; ARTICLE 29 DATA PROTECTION WORKING PARTY, OPINION 05/2014 ON ANONYMISATION TECHNIQUES, 5-6 (2014) [hereinafter *Working Party Opinion on Anonymisation Techniques*].

dates back more than a century and has been known to the general legal academy for half of that time.[7] The "hybrid concept" approach we describe shares with economic analysis of law the application of formal (mathematical) modes of analysis to the design of law, and the general conviction that formal approaches are valuable tools for ensuring internal consistency and avoiding general impossibilities. However, the approach taken by hybrid legal-technical concept development diverges from economic analysis of law in two significant ways. First, it expands the formal theory employed to include computer science and information theory as well as economics. Second, it further emphasizes a practice of conceptualization via a two-way transfer that is facilitated by a multidisciplinary dialog.

Over approximately the last half-century, the field of science and technology studies (which draws from and intersects with sociology, political science, and anthropology) has developed an extensive literature on how scientific ideas are communicated to, integrated in, and regulated by societies. Within this large field, the concept of a "boundary object"[8] may have the closest correspondence to "hybrid concepts."[9] The hybrid concept approach is similar in spirit to that of boundary objects, as both are motivated by the insight that progress across domains requires intellectual artifacts designed to facilitate collaboration. However, whereas the boundary object analysis focuses on the development of tangible and singular artifacts such as a specimen catalog, a physical scale model, or a project requirement document that provides reference descriptions in both domains,[10] hybrid concepts focus on the development of abstract definitions that enable joint problem analysis in both domains.

Finally, in the last ten to twenty years a vast literature has developed on the regulation of new classes of information technologies,[11] complemented by a less

---

[7] For a review of the economic analysis of law, see Lewis Kornhauser, *The Economic Analysis of Law*, STANFORD ENCYCLOPEDIA OF PHILOSOPHY (July 17, 2017), https://plato.stanford.edu/entries/legal-econanalysis/ [https://perma.cc/BB2D-5HSB].

[8] *See* Susan Leigh Star & James R. Griesemer, *Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology*, 19 SOC. STUD. OF SCI. 387, 409-10 (1989).

[9] A boundary object is a specific instance of a physical or informational object (such as a catalog of specimens, or an architectural drawing) that is used to communicate across a technical and a user community – and that is both flexible enough to adapt to local needs and constraints of the several parties using them, yet robust enough to maintain a common identity across communities of users. First used to describe practices in zoology, the concept of boundary objects has since been incorporated into the study of engineering practices, software development, and scientific policymaking, among other areas. For a review of applications of boundary objects, see Pascale Trompette & Dominique Vinck, *Revisiting the Notion of Boundary Object*, 3 REVUE D'ANTHROPOLOGIE DES CONNAISSANCES 3 (2009).

[10] *Id.* at 3-4.

[11] For a notable example, see JONATHAN ZITTRAIN, THE FUTURE OF THE INTERNET—AND HOW TO STOP IT 196-97 (2008). For a current introduction to this area, see IAN LLOYD, INFORMATION TECHNOLOGY LAW (9th ed. 2020).

broad but equally deep literature sparked primarily by Lawrence Lessig's commentary on the ways in which the design of technological infrastructure can act as a form of regulation.[12] Although this article is not intended as a panacea for regulation, or as a commentary on other approaches, we argue that hybrid concepts are a potentially useful tool for regulating information technologies in complex environments. Further, we conjecture that explicit incorporation of hybrid concepts can add nuance to the analysis of "code is law."

II. THE CONCEPT OF SINGLING OUT

As a case study for a hybrid legal-technical approach to privacy analysis, this article explores an anonymization-related concept appearing in Recital 26 of the GDPR called singling out.[13] Section II.A introduces the term ANONYMIZATION and distinguishes it from other possible meanings of the word anonymization. Section II.B analyzes the definitions of anonymization and singling out based on the text of the GDPR and the opinions of the Article 29 Data Protection Working Party. Section II.C identifies and rebuts a simplistic alternative interpretation of singling out as isolation.

A. Terminology: What is ANONYMIZATION?

The term "anonymization" (similarly, "de-identification") is used to refer to a collection of overlapping but fundamentally different concepts. The meanings are very rarely distinguished and are often conflated, muddying the privacy discourse.[14] In one interpretation, anonymizing data means transforming it using a combination of specific techniques, including aggregation, suppression, random swapping, and pseudonymization. In another interpretation, anonymizing data means transforming it in a way so as to guarantee some specific property of the output. A prominent example of this approach is k-anonymity, where a k-anonymization algorithm suppresses and generalizes attributes which are considered

[12] *See* LAWRENCE LESSIG, CODE AND OTHER LAWS OF CYBERSPACE ix-xi (1999).

[13] *See* GDPR *supra* note 6, at Recital 26.

[14] For a discussion of various definitions of "anonymization" and "de-identification"—and how the terms are, in some contexts, used interchangeably or, in others, explicitly distinguished, see NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY, INTERNAL REPORT 8053, DE-IDENTIFICATION OF PERSONAL INFORMATION (2015). *See also* International Standard ISO/IEC 20889, Privacy enhancing data de-identification terminology and classification of techniques 31 (2018) (noting that the standard "does not use the term 'anonymize' . . . because the term has been used in the past to convey a range of different meanings" and providing a table illustrating the relationships between "anonymization," "anonymisation," "de-identification," "pseudonymization," and related terms used in various guidance and standards documents); Jules Polonetsky, Omer Tene & Kelsey Finch, *Shades of Gray: Seeing the Full Spectrum of Practical Data De-identification*, 56 SANTA CLARA L. REV. 593, 596 (2016) (observing that "[a]lthough academics, regulators, and other stakeholders have sought for years to establish common standards for de-identification, they have so far failed to adopt even a common terminology").

potentially identifying to the point that each record in the resulting dataset is identical with at least $k-1$ other records.[15] Another common example is the threshold rule of three, where counts of three or less are zeroed in statistical tables.[16] Under yet another interpretation, anonymizing data means transforming it in a way that makes certain privacy failures—such as re-identification or attribute inference—unlikely or impossible. A fourth interpretation is that anonymizing data means transforming it in a way that frees it from regulation. Under this interpretation, a regulation thus implicitly defines anonymization as that which suffices to do so.

In this article, ANONYMIZATION is used to mean transforming personal data in a way that frees it from GDPR regulation (or a different regulation or statute when specified). The word is stylized to distinguish this particular meaning from other uses of the term in the vernacular or in the data protection literature. Using this styling convention, ANONYMIZING personal data makes it ANONYMOUS, and the act of ANONYMIZING is ANONYMIZATION. Purporting or attempting to ANONYMIZE personal data by means of applying one or more anonymization techniques is not the same as ANONYMIZING it.

It is important to note that to say that personal data has been ANONYMIZED is to make a legal claim: that it is not personal data for the purposes of GDPR regulation. Thus, data which has been pseudonymized, or aggregated, or k-anonymized, or which was analyzed with a differentially private analysis has not necessarily been ANONYMIZED unless it is demonstrated that the technique applied to the data is sufficient to ensure it falls outside of the scope of data protection rules. So too with data that a controller claims to have de-identified or anonymized. Unless the controller makes clear that the anonymization procedure they applied suffices for rendering the data ANONYMIZED – i.e., free from GDPR regulation – it should be assumed that they are not distinguishing the various senses of anonymization.

Although many companies are currently developing and marketing tools for anonymization and de-identification, it is not clear whether such tools should be considered sufficient to demonstrate compliance with existing legal standards for privacy protection, and in particular that they ANONYMIZE data.[17] This

[15] *See* Latanya Sweeney, *k-Anonymity: A Model for Protecting Privacy*, 10 INT'L J. ON UNCERTAINTY, FUZZINESS AND KNOWLEDGE-BASED SYS. 557 (2002).

[16] *See* FEDERAL COMMITTEE ON STATISTICAL METHODOLOGY, STATISTICAL POLICY WORKING PAPER 22 (Second version, 2005), Report on Statistical Disclosure Limitation Methodology (2005).

[17] *See, e.g.*, Jules Polonetsky et al., *Shades of Gray: Seeing the Full Spectrum of Practical Data De-identification*, 56 SANTA CLARA L. REV. 593, 596 (2016) ("Despite a broad consensus around the need for and value of de-identification, the debate as to whether and when data can be said to be truly de-identified has appeared interminable."); Paul Francis, *The 4 Pillars of GDPR Compliance Assurance in Data Anonymization*, AIRCLOAK, (Feb. 11, 2019), https://aircloak.com/gdpr-compliance-assurance/ [https://perma.cc/6RPG-RNC3] (observing that, although "[t]he GDPR also requires certification programs, including for certifying anonymity, . . . to date no organizations have done this for anonymity, nor to my knowledge

uncertainty stems in large part from the parallel development of regulatory and technical approaches to data protection and their conceptual differences.

What does it take for anonymization to be ANONYMIZATION? What is required of a transformation to be considered ANONYMIZATION, and which transformations satisfy the requirements? These questions motivate this work.

B. ANONYMIZATION and singling out in the GDPR

The GDPR regulates the processing of personal data, and its material scope is delineated in Article 1: "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."[18] The GDPR, therefore, places no restrictions on the processing of data which is not personal, including personal data that has been rendered ANONYMOUS, as 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."[19]

The concept of personal data defined by the GDPR is broad. Article 4 defines personal data—by definition, non-ANONYMOUS—as "any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly."[20] What it means for a natural person to be "identified, directly or indirectly"[21] is further elaborated in Recital 26: "To determine whether a natural person is identifiable, account should be taken of all 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."[22] Recital 26 clarifies that in "means reasonably

does any have a clear plan on how to do so" and that "[t]his creates a problem for the providers and users of anonymization technologies").

[18] GDPR *supra* note 6, at Article 1.

[19] *Id*. at Recital 26.

[20] *Id*. at Article 4.

[21] We do not address the meaning of "information relating to [a . . .] natural person" in this article. Other scholars have argued that the "identified, directly or indirectly" criterion is the element most in need of a close analysis, as the "natural person" criterion "does not involve any significant critical discussion once it is clarified that data protection law applies only to natural persons as opposed to legal persons" and "the criterion of data 'relating to' the data subject . . . can, in most cases be subsumed in the criterion of 'identifiability' in a sense that data that is capable of identifying a person can be assumed to also relate to the data subject." Worku Gedefa Urgessa, *The Protective Capacity of the Criterion of 'Identifiability' under EU Data Protection Law*, 2 EUR. DATA PROT. L. REV. 521, 521 n.1 (2016) (citing LEE BYGRAVE, DATA PRIVACY LAW: AN INTERNATIONAL PERSPECTIVE 130–131 (2014)).

[22] GDPR *supra* note 6, at Recital 26 (emphasis added). It is instructive to compare the text from Recital 26 of the GDPR with that of the Data Protection Directive: "[. . .] to determine whether a person is identifiable, account should be taken of all the means likely reasonably to be used either by the controller or by any other person to identify the said person." *See*

likely to be used" the standard refers to attacks which are feasible in terms of technology and cost.[23]

The phrase "singling out" appears nowhere else in the regulation, nor did it appear in the GDPR's predecessor, the 1995 Data Protection Directive (DPD).[24] One must look elsewhere to understand the meaning of singling out and its relationship to ANONYMIZATION.

This Section explores singling out in the context of processing information about natural persons, proceeding in three parts. First, it discusses guidance from the most persuasive authority available for understanding singling out as used in the GDPR, i.e., the opinions of the Article 29 Data Protection Working Party (henceforth, the Working Party).[25] Second, as used by the Working Party, to single out is to specify a collection of attributes that distinguishes an individual from all other individuals included in the data.[26] We call this notion of singling out *isolation*.[27] Third, it concludes that ANONYMIZATION under the GDPR requires preventing singling out. Equivalently, if singling out is possible, then the dataset is personal data.

Before continuing, it is worth addressing some common misconceptions about ANONYMIZATION and singling out. Whether or not data has been ANONYMIZED depends on the *relationship between the original personal data and released data*, not on the released data alone.[28] This is because it is

Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. While the DPD does not mention any of the "means likely reasonably to be used" the GDPR adds mention of singling out but does not specify other means.

[23] *See* GDPR, *supra* note 6, at Recital 26 ("To ascertain whether means are reasonably likely to be used to identify the natural person, account should be taken of all objective factors, such as the costs of and the amount of time required for identification, taking into consideration the available technology at the time of the processing and technological developments.").

[24] *See* Directive 95/46/EC, *supra* note 22.

[25] For a discussion explaining why this guidance is the most persuasive authority on the meaning of singling out even though the Working Party has since been replaced by the European Data Protection Board, see *infra* Subsection II.B.1.

[26] This interpretation was originally proposed without analysis in Paul Francis, Sebastian Probst-Eide, Pawel Obrok, Cristian Berneanu, Sasa Jurie & Reinhard Muz, *Extended Diffix*, Working Paper 3 (2018), https://aircloak.com/wp-content/uploads/Complete-Diffix.pdf [https://perma.cc/CX75-S87V] ("We define singling out as occurring when an analyst correctly makes a statement of the form 'There is exactly one user that has these attributes.'").

[27] Section II.C and onward argue that isolation is an incomplete theory of singling out.

[28] For an intuitive explanation of why privacy is properly understood as a property of the relationship between the input and output of a computation rather than a property of the output alone, see Alexandra Wood, Micah Altman, Aaron Bembenek, Mark Bun, Marco Gaboardi, James Honaker, Kobbi Nissim, David R. O'Brien, Thomas Steinke & Salil Vadhan, *Differential Privacy: A Primer for a Non-technical Audience*, 21 VAND. J. ENT. & TECH. L. 209, 221–225 (2018). The interpretation of ANONYMIZATION in this sentence is arguably reflected by the terminology used in Recital 26 and guidance from the Working Party. *See* GDPR, *supra* note 6, at Recital 26 (referring to "personal data rendered anonymous in such a manner

impossible to fully assess disclosure risks by inspection of the data release in isolation.[29] For instance, consider a data release that states that "Fatima Portokalos is HIV-positive." Looking at the data release alone and not knowing how it was created, it impossible to determine whether it relates to an identifiable natural person, or whether instead it was created by making up a fake name and assigning it an arbitrary disease.[30] Accordingly, this article examines the protections offered by *mechanisms*, or processes, that aim to transform personal data into data releases that are simultaneously non-disclosive and useful.

Orienting one's conception of ANONYMIZATION from the data release to the mechanism that creates it helps to dispel another misconception: that the *format* of the data release matters. While often used in reference to mechanisms that output microdata comprised of records corresponding to distinct individuals, the concepts of ANONYMIZATION and singling out apply to other forms of data release, including models created using machine learning on personal data.[31] What matters is not in what form the data occurs, but what can be done with it.

Singling out a person within a dataset should not be conflated with learning the person's name.[32] A person may be singled out even if the data is pseudonymous and the person's name is very difficult to learn, for example, by specifying a combination of attributes that is unique in the dataset. Moreover, a common name or surname may not be sufficient to single an individual out from a large population (e.g., a country) though it may suffice to single out within a smaller population (e.g., a classroom).[33] However, singling out – even without a name – can be seen as a stepping stone to a greater privacy breach. For example, singling

---

that the data subject is not or no longer identifiable"); *see also Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 6 (using the term "anonymisation technique" instead of "anonymity" or "anonymous data").

[29] *See* Wood et al., *supra* note 28, at 221–225.

[30] *See id.*

[31] *See Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 12 ("A dataset is a collection of records that can be shaped alternatively as a table (or a set of tables) or as an annotated/weighted graph, which is increasingly the case today. The examples in this opinion will relate to tables, but they are applicable also to other graphical representations of records."); s*ee also* Michael Veale, Reuben Binns & Lilian Edwards, *Algorithms that remember: Model inversion attacks and data protection law*, 376 PHIL. TRANSACTIONS ROYAL SOC'Y A 20180083 (2018) (analyzing CJEU opinions and Article 29 Working Party guidance and concluding that "information on an individual's membership of a training set would indeed fall within the scope of personal data, regardless of how trivial or mundane it might be to the individual it concerns" and that "model inversion and membership inference attacks, where possible, do risk models being considered as personal data even without resorting to a maximalist reading of data protection law").

[32] *See Working Party Opinion on Personal Data*, *supra* note 6, at 14 ("[W]hile identification through the name is the most common occurrence in practice, a name may itself not be necessary in all cases to identify an individual. This may happen when other 'identifiers' are used to single someone out.") (emphasis in original).

[33] *See id.* at 13.

out a person using a unique combination of attributes may enable a linkage between the person's record in the dataset to some external source of information. This is how the linkage between users of Netflix and the Internet Movie Database was achieved in a well-known de-anonymization demonstration.[34]

1. Guidance from the Article 29 Data Protection Working Party

Singling out appears multiple times in guidance from the Article 29 Data Protection Working Party, an advisory body set up under Article 29 of the 1995 Data Protection Directive.[35] Most notably, the Working Party hints at the definition of singling out in a 2007 opinion on the concept of personal data and more fully develops it in its 2014 opinion on anonymization techniques, establishing singling out as one of three criteria for effective anonymization.[36] We believe that, as of the date of this article, these two Working Party opinions, taken together, are the most persuasive authority on the meaning of singling out in the GDPR.[37] This conclusion is supported by the following. First, the interpretation of singling out set forth in these two opinions has recently been adopted by many EU data protection authorities in their guidance on ANONYMIZATION.[38] Second,

[34] *See* Arvind Narayanan & Vitaly Shmatikov, *Robust De-anonymization of Large Sparse Datasets*, 29 PROC. IEEE SYMP. ON SEC. AND PRIV. 111 (2008).

[35] The concept of singling out is referenced in Working Party Opinion 04/2007, Opinion 01/12, Working Document 02/2013, and Opinion 05/2014, among others. *See Working Party Opinion on Personal Data*, *supra* note 6; Article 29 Working Party, Opinion 01/2012 on the data protection reform proposals (2012); Article 29 Working Party, Working Document 02/2013 providing guidance on obtaining consent for cookies (2013); *Working Party Opinion on Anonymisation Techniques*, *supra* note 6.

[36] *See Working Party Opinion on Personal Data*, *supra* note 6; *Working Party Opinion on Anonymisation Techniques*, *supra* note 6.

[37] With the implementation of the GDPR, the Article 29 Data Protection Working Party has been replaced by the European Data Protection Board (EDPB). In May 2018, the EDPB acknowledged "the continuity of the work provided by the predecessor Article 29 Working Party" and endorsed some of its documents "[w]ithout prejudice to any future revision as appropriate." *See* European Data Protection Board, Endorsement 1/2018 (May 25, 2018), https://edpb.europa.eu/sites/edpb/files/files/news/endorsement_of_wp29_documents.pdf [https://perma.cc/MN7Z-EJBH]. While the EDPB has not endorsed Opinion 04/2007, Opinion 05/2014, or any of the other Working Party documents that mention singling out, these opinions remain an influential source of interpretive guidance on the concepts of personal data, anonymization, and singling out as applied in EU data protection law. *See, e.g.*, Letter from Wojciech Rafał Wiewiórowski, European Data Protection Supervisor, to Roberto Viola, Directorate General of Communication, Networks, Content and Technology at the European Commission (Mar. 25, 2020) (citing *Working Party Opinion on Anonymisation Techniques*), https://edps.europa.eu/sites/edp/files/publication/20-03-25_edps_comments_concerning_covid-19_monitoring_of_spread_en.pdf [https://perma.cc/TNH9-Q6LQ].

[38] Many data protection authorities have adopted the Working Party's view that ANONYMIZATION requires preventing singling out, linkability, and inference, as well as definitions of singling out that closely track the language from the Working Party opinion on anonymization techniques. *See, e.g.,* CNIL [France], *Sheet n°1: Identify personal data* (June

during the data protection reform process, the Working Party urged the European Commission to ensure the GDPR's definition of personal data takes into account this conception of the risk of "singling out."[39] Third, early drafts of Recital 26 "seemingly adopted" the approach to singling out set forth in the Working Party opinion on anonymisation techniques.[40] Fourth, no alternative definition has been presented by an authority such as the CJEU or EDPB. Fifth, the definitions of personal data, to which singling out is inextricably linked,[41] under the GDPR and Data Protection Directive are substantially similar.[42]

11, 2020), https://www.cnil.fr/en/sheet-ndeg1-identify-personal-data [https://perma.cc/HC2J-F62J]; DATA PROT. COMM'N [IRELAND], GUIDANCE NOTE: GUIDANCE ON ANONYMISATION AND PSEUDONYMISATION 3, 5-6 (June 2019).

[39] *See* Opinion 01/2012, *supra* note 35, at 9 ("A natural person can be considered identifiable when, within a group of persons, he or she can be distinguished from the other members of the group and consequently be treated differently. This has been set out in the earlier adopted opinion of the Working Party on the concept of personal data (WP136). Recital 23 should therefore be amended in order to clarify that the notion of identifiability also includes singling out in this way."); Article 29 Data Protection Working Party, Opinion 08/2012 providing further input on the data protection reform discussions, at 5 (2012) ("One of the main conclusions of this analysis [in the opinion on the concept of personal data] is that a natural person can be considered identifiable when, within a group of persons, he or she can be distinguished from other members of the group and consequently be treated differently. It is therefore suggested to clarify in Recital 23 and Article 4 that the notion of identifiability also includes singling out in this way."); Letter from the Article 29 Data Protection Working Party on Trilogue to Věra Jourová, Commissioner for Justice, Consumers and Gender Equality of the European Commission 2, 2 (June 17, 2015), https://ec.europa.eu/justice/article-29/documentation/other-document/files/2015/20150617_letter_from_the_art29_wp_on_trilogue_to_msjourova_en.pdf [https://perma.cc/C76M-3TWN] (arguing that "[t]o ensure the general objective of maintaining a high-level of protection of personal data is upheld, personal data should be defined in a broad manner in line with technological evolution" and that "the definition of personal data should therefore take into account the situation in which people can be 'singled out' on the basis of identifiers or other information and could subsequently be treated differently").

[40] *See* Leslie Stevens, *The Proposed Data Protection Regulation and Its Potential Impact on Social Sciences Research in the UK*, 2 EUR. DATA PROTECTION L. REV. 97, 104 (2015) ("The Article 29 Working Party's more recent approach to 'singling out' [set forth in Opinion 05/2014 on Anonymization Techniques] was seemingly adopted into the wording of Parliament, and less explicitly in the Council's text for Recital 23, representing a departure from the more proportionate approach offered under the DPD and previously by the Working Party itself [in Working Document 02/2013]."). Note that, although the final text of Recital 26 differs from that of the Parliament draft, the final text remains consistent with the Working Party Opinion on Anonymization Techniques in explicitly recognizing singling out as one means of identification.

[41] *See infra* Section II.B.2.

[42] *See* GDPR, *supra* note 6, at Article 4 ("'personal data' means any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more

2. What is the definition of singling out?

The Working Party's 2007 opinion on the concept of personal data and 2014 opinion on anonymization techniques strongly suggest that the Working Party equated singling out with *isolation*: the act of specifying a collection of attributes that distinguishes an individual from all other individuals included in a given dataset.[43] The remaining opinions offer only very limited clues for or against singling out as isolation, mentioning singling out only in passing or in reference to the above two opinions.[44]

The opinion on anonymization techniques defines singling out as "the possibility to isolate some or all records which identify an individual in the dataset."[45] A more coherent reading is that to single out is *to isolate* some or all records *corresponding to* an individual in the dataset. This alternative resolves three problems with the original and generally makes the opinion more cogent. First, the original definition suffers from a grammatical mismatch. "Singling out" is a verb (in present participle form);[46] a "possibility" of isolating is a noun. Second, the original definition is incongruous with the use of singling out throughout the opinion as something that can "possibly" be done with or to data (i.e., a possibility of a possibility of isolating).[47] Third, the definition's use of "identify" introduces a conceptual circularity: singling out makes a person identifiable as a legal matter,[48] but is itself defined with reference to identifying a person. Elsewhere the opinion reinforces the former point,[49] and challenges the latter. It is

factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person"); Council Directive 95/46/EC, art. 2, 1995 O.J. (L. 281) ("'personal data' shall mean any information relating to an identified or identifiable natural person ('data subject'); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity").

[43] *See Working Party Opinion on Personal Data*, *supra* note 6; *see also Working Party Opinion on Anonymisation Techniques*, *supra* note 6.

[44] The concept of singling out is referenced in Working Party Opinion 04/2007, Opinion 01/12, Working Document 02/2013, and Opinion 05/2014. *See Working Party Opinion on Personal Data*, *supra* note 6; Opinion 01/2012, *supra* note 35; Working Document 02/2013, *supra* note 35; *Working Party Opinion on Anonymisation Techniques*, *supra* note 6.

[45] *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 11; *see* Opinion 09/2014 on the application of Directive 2002/58/EC to device fingerprinting 4, n.10 (Nov. 25, 2014).

[46] *See Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 3, 9, 10. (stating "to single out," "prevents all parties from singling out," and "allow an individual data subject to be singled out").

[47] *Id.* at 13, 14, 21.

[48] *See infra* Subsection II.B.3.

[49] *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 9 ("An effective anonymisation solution prevents all parties from singling out an individual in a dataset"); *id.* at 10 ("It must be clear that 'identification' not only means the possibility of retrieving a

the records "of an individual" that get singled out, not records that "identify an individual."[50] We use "corresponding to an individual" to emphasize that it is a record's relationship to the individual that matters, not any sort of ownership.

It remains to understand what it means to isolate records which correspond to an individual in a dataset. The opinion provides numerous illustrations, including three which follow a common pattern:

Example 1: Singling out by fingerprinting

The Working Party opinion on anonymization techniques provides three examples of singling out in de-identified datasets using a "fingerprint" that uniquely describes a single individual in the dataset. All are described as failures of anonymization that make them vulnerable to re-identification, even though only the Netflix example describes the actual re-association of identities with the pseudonymized records.

The examples illustrate that the essence of GDPR singling-out is what we call *isolation*: the act of specifying a collection of attributes that distinguishes an individual from all other individuals included in a given dataset.

- "Researchers at MIT recently analyzed a pseudonymised dataset consisting of 15 months of spatial-temporal mobility coordinates of 1.5 million people on a territory within a radius of 100 km. They showed that 95% of the population could be singled-out with four location points, and that just two points were enough to single-out more than 50% of the data subjects[.]"[51]

- "A very famous re-identification experiment is the one performed on the customers' database of the video content provider Netflix. . . . In spite of [being 'anonymised'], it was found that 99% of user records could be uniquely identified in the dataset using 8 ratings and dates with 14-day errors as selection criteria, whilst lowering the seletion criteria (2 ratings and 3-day error) still allowed identifiying 68% of users."[52]

person's name and/or address, but also includes potential identifiability by singling out, linkability and inference.").

[50] *Id.* at 13 ("It is still possible to single out the records of an individual (perhaps in a non-identifiable manner)"); *see also id.* at 14, 21.

[51] *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 23 (citing Yves-Alexandre de Montjoye, César A. Hidalgo, Michel Verleysen & Vincent D. Blondel, *Unique in the Crowd: The Privacy Bounds of Human Mobility*, 3 NATURE SCI. REP. 1376 (2013)).

[52] *Id.* at 13 (citing Arvind Narayanan & Vitaly Shmatikov, *Robust De-anonymization of Large Sparse Datasets*, PROC. OF THE 2008 IEEE SYMP. ON RSCH. IN SEC. AND PRIV. 111 (2008)).

"The auxiliary information retrieved in the IMDB database could be imported into the released Netflix data set, thus enriching with identities all the supposedly anonymised records."[53]

- "[A]n identification experiment performed against a social network[] exploited the social graph of users pseudonymised by means of labels. In this case, the attributes used for identification were the list of contacts of each user as it was shown that the likelihood of an identical list of contacts between two individuals is very low. Based on this intuitive assumption, it has been found that a sub-graph of the internal connections of a very limited number of nodes constitutes a topological fingerprint to retrieve, hidden within the network, and that a wide portion of the entire social network can be identified once this sub-network has been identified."[54]

The examples help us understand what it means to isolate records corresponding to an individual in a dataset. In each example, the act of the implicit hypothetical attacker is the act of specifying a collection of attributes: four location points, eight ratings and dates, the overlap between a user's social graph and a certain small sub-network. These attributes distinguish an individual within the data: one and only one person matches the specified attributes.

This interpretation of singling out as isolation is supported by the Working Party opinion on the concept of personal data. That opinion makes three relevant references to singling out. Each refers to an act of specifying attributes or information that "allow the individual to be distinguished from others."[55] A later Working Party opinion – providing feedback on proposed language for what would later become Recital 26 – paraphrases:

> A natural person can be considered identifiable when, within a group of persons, he or she can be distinguished from the other members of the group and consequently be treated differently. This has been set out in the earlier adopted opinion of the Working Party on the concept of personal data (WP136). Recital 23 should therefore be amended in order to clarify that the notion of identifiability also includes singling out in this way.[56]

There is very little scholarship about the definition of singling out per se. Paul Francis et al. originally proposed a version of the singling-out-as-isolation view,

[53] *Id.* at 30.

[54] *Id.* at 30 (citing Lars Backstrom, Cynthia Dwork & Jon Kleinberg, *Wherefore Art Thou R3579X? Anonymized Social Networks, Hidden Patterns, and Structural Steganography*, PROC. 16TH INT'L CONF. ON WORLD WIDE WEB 181 (2007)).

[55] *Working Party Opinion on Personal Data*, *supra* note 6, at 13–14.

[56] Opinion 01/2012, *supra* note 35, at 9. *But see* Leslie Stevens, *The Proposed Data Protection Regulation and Its Potential Impact on Social Sciences Research in the UK*, *supra* note 40, at 104 (2015) (discussed *infra*).

albeit without analysis.[57] The ISO/IEC standard on data de-identification terminology adopts a similar definition of singling out, and refers to the Working Party opinion on anonymization techniques to support its statement that "[p]seudonymization used alone does not reduce the risk that an individual data principal can be singled out."[58] Leslie Stevens argues that the excerpt quoted above supports an impact-oriented view of singling out.[59] She argues that "the Working Party is concerned with singling out when it would result in *treating individuals differently*, when 'singling out' would result in impact upon individuals."[60] However, Stevens ultimately concludes that the later opinion on anonymization techniques rejects this view and that the later approach was "seemingly adopted" during the trilogue negotiations (wherein the European Commission, European Parliament, and Council of the European Union reconcile differences in proposed legislation).[61]

3. What is the legal significance of singling out?

The GDPR recognizes singling out as one way to identify a person in data: if, taking account of all the means reasonably likely to be used, the data allows a person to be singled out, then that data is personal data.[62] This understanding of the relationship between singling out and the definition of personal data is substantiated by the text of Article 4 and Recitals 26 and 28.[63] As Damian Clifford and Jef Ausloos explain:

---

[57] Francis et al., *supra* note 26, at 3 (emphasis omitted) ("We define singling out as occurring when an analyst correctly makes a statement of the form 'There is exactly one user that has these attributes.'").

[58] ISO/IEC 20889:2018 at 5 (defining "single out" as "isolate records belonging to a data principal in the dataset by observing a set of characteristics known to uniquely identify this data principal"); *see also id.* at 15. Note that under the ISO definition, singling out is an act of "observing" certain attributes (rather than "specifying") and requires that the attributes be "known to" distinguish an individual from all other individuals (rather than merely having that effect).

[59] Stevens, *supra* note 40, at 104 (citing Article 29 Data Protection Working Party, Working Document 02/2013 providing guidance on obtaining consent for cookies and *Working Party Opinion on Anonymisation Techniques*, *supra* note 6).

[60] Stevens, *supra* note 40, at 104 (emphasis in original).

[61] *Id.*

[62] *See* ROSEMARY JAY, GUIDE TO THE GENERAL DATA PROTECTION REGULATION: A COMPANION TO DATA PROTECTION LAW AND PRACTICE 339 (2017) ("It is clear that as long as a person can be 'singled out' they are regarded as identifiable.").

[63] Note that some scholars have argued in favor of an alternative interpretation of this relationship. *See* Miranda Mourby, Elaine Mackey, Mark Elliot, Heather Gowans, Susan E. Wallace, Jessica Bell, Hannah Smith, Stergios Aidinlis & Jane Kaye, *Are 'pseudonymised' data always personal data? Implications of the GDPR for administrative data research in the UK*, 34 COMPUTER L. & SECURITY REV. 222, 228 (2018) ("Singling out need only be considered if it is a means 'reasonably likely' to be used to [identify a person]"). As we explain in this section, we believe this alternative interpretation is contradicted by both the text of the

> The definition of personal data combined with the definition of pseudonymisation (Article 4(5) GDPR) and the clarification regarding the interaction between these two definitions (provided in Recitals 26 and 28), indicate that any data capable of "singling out" an individual should be considered as personal. This clarification should be recognised as a significant take-away given that a failure to include such data within the scope of the definition of personal data would have undermined the protections provided by the framework.[64]

In other words, ANONYMIZING personal data requires preventing singling out, and, in order to fall outside the scope of the regulation, data must not allow singling out. Note, however, that the converse does not necessarily hold because there may be other means reasonably likely to be used to identify a person. Indeed, relevant guidance—such as that released by the Working Party—considers means of identifying such as linkability and inference in addition to singling out.[65] Though it is only one facet of ANONYMIZATION, singling out provides a natural test: if a technique allows singling out (when taking account of all the means reasonably likely to be used), then it does not ANONYMIZE under GDPR.

Many companies claim to achieve ANONYMIZATION, whether to sell ANONYMIZATION-as-a-service[66] or to enable secondary uses of data.[67] For example, telecommunications operators, including Orange in France and Deutsche Telekom in Germany, have begun sharing purportedly anonymized mobile phone users' location data with EU governments and public health organizations for the purpose of tracking mobility patterns during the COVID-19 pandemic with the understanding that ANONYMIZED data falls outside the scope of the EU's data protection rules.[68] Related arguments are also being used to justify the use

GDPR (Article 4, Recital 26, and Recital 28) and the guidance from the Article 29 Data Protection Working Party documented in the Opinion 05/2014 on Anonymisation Techniques. GDPR, *supra* note 6, at Article 4, Recital 26, Recital 28; *Working Party Opinion on Anonymisation Techniques, supra* note 6.

[64] Damian Clifford & Jef Ausloos, *Data Protection and the Role of Fairness* 6-7 (KU Leuven Ctr. for IT & IP Law, CiTiP Working Paper 29/2017, 2017) (citation omitted).

[65] *See Working Party Opinion on Anonymisation Techniques*, *supra* note 6.

[66] *See, e.g.*, AIRCLOAK, https://aircloak.com [perma.cc/YZ4X-WMSG]; PRIVITAR, https://www.privitar.com [perma.cc/RY4Y-MJA8]; ANONOS, https://www.anonos.com [perma.cc/93CD-B3LD].

[67] *See, e.g.*, Paige Maas, Shankar Iyer, Andreas Gros, Wonhee Park, Laura McGorman, Chaya Nayak & Alex Dow, *Facebook Disaster Maps: Aggregate Insights for Crisis Response & Recovery*, PROC. 16TH ISCRAM CONF. 836, 837 (2019).

[68] *See* Mark Scott, Laurens Cerulus & Laura Kayali, *Commission tells carriers to hand over mobile data in coronavirus fight*, POLITICO (Mar. 23, 2020, 11:23 PM), https://www.politico.eu/article/european-commission-mobile-phone-data-thierry-breton-coronavirus-covid19 [https://perma.cc/BZ7X-SYUF]; Natasha Lomas, *Telco metadata grab is for modelling COVID-19 spread, not tracking citizens, says EC*, TECHCRUNCH (Mar. 27, 2020, 12:05 PM), https://techcrunch.com/2020/03/27/telco-metadata-grab-is-for-modelling-covid-19-spread-not-tracking-citizens-says-ec [https://perma.cc/FL5X-CG4L].

of secure multiparty computation to perform computations on personal data.[69] An analysis of their vulnerability to singling out can be used to evaluate these claims. The question of evaluating claims of ANONYMIZATION is not new. Some scholars question whether ANONYMIZATION is possible at all.[70] Others propose risk-based approaches to ANONYMIZATION and risk assessment frameworks for evaluating ANONYMIZATION tools.[71] The Working Party and the data protection authorities of many member states have published technical guidance documents on assessing the effectiveness of ANONYMIZATION techniques.[72]

C. Isolation is an incomplete theory of singling out

In the preceding Section, we argued that the most persuasive theory of singling out in the GDPR, as reflected in the Working Party opinions, equates singling out with what we call isolation. This Section demonstrates that the theory is incomplete.

Isolation is the act of specifying a collection of attributes that distinguishes an individual from all other individuals included in a given dataset. Francis et al. propose an operative version of this view that singling out is isolation, as described in Example 2.[73]

[69] *See, e.g.*, *Cybernetica*, *Sharemind SDK 2019.03 Beta*, https://sharemind-sdk.github.io [https://perma.cc/P876-2MNV] ("Sharemind is a novel database and application server that collects data in an encrypted form and uses techniques like homomorphic encryption, secure multi-party computation and hardware isolation [...]. As such, Sharemind helps achieve compliance with the General Data Protection Regulation (GDPR) and similar laws.").

[70] *See, e.g.*, Paul Ohm, *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*, 57 UCLA L. REV. 1701, 1703 (2010).

[71] *See, e.g.*, MARK ELLIOT, ELAINE MACKEY, KIERON O'HARA & CAROLINE TUDOR, THE ANONYMISATION DECISION-MAKING FRAMEWORK vi (2016).

[72] *Working Party Opinion on Personal Data*, *supra* note 6, at 18; *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 3-4. *See also, e.g.*, INFORMATION COMMISSIONER'S OFFICE, ANONYMISATION: MANAGING DATA PROTECTION RISK, CODE OF PRACTICE, 2012, 18-27 (UK).

[73] Francis et al., *supra* note 26 (emphasis omitted).

Example 2: Singling out as isolation (Francis et al.)

"We define singling out as occurring when an analyst correctly makes a statement of the form 'There is exactly one user that has these attributes. [An] analyst may claim that there is a single user with attributes [gender = 'male', age = 48, zipcode = 48828, lastname = 'Ng']. If this is true, then the analyst has correctly singled out that user. The attributes don't need to be personal attributes as in this example. If the analyst correctly claims that there is a single person with the geo-location attributes [long = 44.4401, lat = 7.7491, time = '2016-11-28 17:14:22'], then that person is singled out."

Preventing isolation requires that it be impossible for an analyst to correctly make a statement of the form "There is exactly one user that has such-and-such attributes." Unfortunately, this is mathematically impossible, as illustrated in Example 3.

Example 3: Isolation by random guessing

Consider a university dataset consisting of students' information, including their birthdays (month and day). For simplicity, assume that each birthday is equally likely to be one of the 365 days in a year. If the dataset holds information about 365 students, then, on average one of the 365 students in the dataset will have their birthday on each day of the year. It is, therefore, likely that there is a single person in the dataset with the birthdate attributes [month = 'October', day = 23]. A simple calculation shows that this would be the case with 37% chance, as shown in Technical Details 1 in Subsection III.C.2. Hence, an attacker claiming that there is a single student with attributes [month = 'October', day = 23] would succeed with 37% chance.

The attacker in Example 3 succeeds with 37% chance without consulting the university dataset whatsoever. This means that, regardless of how secure the release of the dataset is, an attacker would succeed in isolating with at least 37% chance. A few notes about the example are in order.[74] First, the attributes [month = 'October', day = 23] were selected arbitrarily; any other date would have the same chance of isolation. Second, this example is an specific instance of a general trend. In typical datasets, a similar type of isolation is always possible with 37% chance. Third, the attacker isolates with a rather high probability—i.e.,

[74] For a more detailed analysis of this example, see discussion *infra* Subsection III.C.2.

roughly one in every three attempts results in a successful isolation.[75] Fourth, having pinpointed a potential target, it is often possible for an attacker to check whether isolation occurred, e.g., by querying the data source or finding corroboration elsewhere. Lastly, the attack is extremely simple and hence meets the "means reasonably likely" standard of Recital 26 in terms of technology and cost.[76]

For these reasons, isolation is an incomplete theory of GDPR singling out. Otherwise, one must conclude that either GDPR ANONYMIZATION is impossible and hence a vacuous concept, or that GDPR ANONYMIZATION allows for a high probability of singling out. Both resolutions seem incongruous with the regulators' intent.

Other scholars have argued that ANONYMIZATION is impossible as a practical matter, that people are simply so distinct and sources of data so numerous that it is folly to consider every possible means of re-identification attack. Example 3 illustrates a very different type of impossibility – mathematical impossibility. There is an inherent risk of isolation, even if it were possible to know everything about every person and to exert total control over all possible data sources. Moreover, the risk remains even if no data is published at all.

III. REFINING AND FORMALIZING THE GDPR'S NOTION OF SINGLING OUT

This Section derives a precise mathematical version of a GDPR singling-out attack – called *predicate singling-out* – which we use to study ANONYMIZATION under the GDPR.[77] The moniker "predicate" distinguishes the mathematical notion from the legal concept of singling out referred to by the GDPR (and, to the extent that they are related, other legal conceptions of distinguishability). Though distinct, we argue that the concepts are related, and that predicate singling-out is germane to a coherent legal concept of singling out, and, by extension, ANONYMIZATION. Namely, security against predicate singling-out attacks is a necessary (but not sufficient) requirement for substantiating a claim that "account [is] taken of all the means reasonably likely to be used [. . .] to identify the natural person directly or indirectly," as is required by Recital 26 of the GDPR.[78]

This Section deserves a roadmap as it is rather long and, in parts, technical. Section III.A is a brief overview of predicate singling-out, introducing the terminology and main ideas underlying the remainder of the article. Section III.B lays out the principles underlying our modeling approach. The remaining three sections (Sections III.C–III.E) are more technical, and readers may choose to

[75] An attacker making three (independent) isolation attempts would succeed at least once with probability 75%. With five attempts, the probability is 90% and with 10 attempts it is 99%. *See* Cohen & Nissim, *supra* note 1, at 8346.

[76] *See supra* note 23.

[77] The mathematical claims are backed by an article previously published by two of the authors of this article. *See* Cohen & Nissim, *supra* note 1, at 8344.

[78] GDPR, *supra* note 6, at Recital 26.

skim these sections during a first read. Section III.C develops and presents our core concept: *predicate singling-out*. It builds on the overview in Section III.A by explaining the many modeling choices and details underlying the definition and providing additional intuition. To illustrate how a mechanism can be analyzed using the definition, Section III.D gives an example of a mechanism that prevents predicate singling-out attacks. Finally, Section III.E discusses the (non-)composability of predicate singling-out and potential implications for the GDPR's notion of singling out.

A. Predicate singling-out in a nutshell

Very roughly, a person is *predicate singled-out* in a dataset if they are described in a way that distinguishes them from all others in the dataset, and the description is so specific as to be unattributable to chance.[79]

A privacy mechanism *prevents predicate singling-out attacks* if the chance that an attacker manages to predicate single-out a person in the dataset is not too different depending on whether the attacker knows the output of the mechanism applied to the dataset, or knows nothing at all about the dataset. We argue that any privacy mechanism that ANONYMIZES data under the GDPR in all cases must prevent predicate singling-out attacks.

Our focus is the privacy mechanism, not a data release. As explained in Section II.B, whether or not data has been ANONYMIZED depends on the *relationship between the original personal data and released data*, not on the released data alone.

Predicates. A *predicate* is a logical statement over variables that may evaluate to either True or False depending on the assignment of values to its variables. In the context of this paper, predicates are used to describe persons. Hence, a predicate is specified by stating some properties of a person's data. The predicate assigns a value True—in which case the predicate is satisfied—or False—in which case the predicate is not satisfied—to a person's data depending on whether or not the person's data agrees with the specified properties. It may be that the data of zero, one, or several persons satisfy a predicate.[80]

For example, the predicate

$$p: \text{[age = 49 AND height } \leq 5\text{'8"] OR [rent } \leq \frac{\text{salary}}{3}\text{]}$$

evaluates to True if at least one of two conditions is met: either a person's age is 49 and their height is at most 5'8", or their rent is at most a third of their salary (or both).

Isolation. A predicate *isolates* a person in a dataset if she is the only person in the dataset that satisfies the predicate. It distinguishes her by returning True on her data and False on everybody else's data.

---

[79] *See* Cohen & Nissim, *supra* note 1, at 8345.

[80] *Id.*

Isolating a person is not sufficient to define predicate singling-out. It is possible to isolate individuals in a dataset purely by chance, even knowing nothing about the dataset. As in Example 3, it turns out that in a dataset of 365 random people, there is about a 37% chance that there is exactly one person born on, say, March 12. However, this is not a meaningful singling-out attack against the dataset because it did not use the dataset in any way.[81]

Baseline risk. Taking into account the probability that isolation would happen by a random guess, we define a statistical *baseline risk* of isolation as the chance than an attacker manages to guess a predicate that isolates a person in a dataset about which they know nothing at all.[82] The baseline depends on how *rare* the predicate is, i.e., how likely the predicate is to output True on a random person in the underlying population.[83] The predicate [born on March 12] returns True on about 1/365 of the population, while the probability that the predicate [born on March 12, and is Jewish, Colombian, vegan, fluent in Mandarin, and a pilot] evaluates to True is very close to 0.

For exceedingly rare predicates like the one above, the baseline risk is also very close to 0.[84] Namely, it is very unlikely that an attacker isolates a person using a rare predicate by chance without seeing the anonymization mechanism's output. Thus, successful isolation can be attributed to the attacker's knowledge of the output. Further, observe that, while a more common predicate might suffice to isolate a person in a dataset, the predicate would be satisfied by the data of many people in the underlying population. The rarer a predicate is, the more it distinguishes individuals in the underlying population.

Predicate singling-out. Putting it all together, a privacy mechanism prevents predicate singling-out attacks if the chance that an attacker manages to isolate a

[81] *See id.* at 8345.

[82] *See id.* at 8348.

[83] *See id*. at 8346.

[84] *See id.*

person in the dataset using an exceedingly rare predicate is small. Otherwise, it enables predicate singling-out attacks.

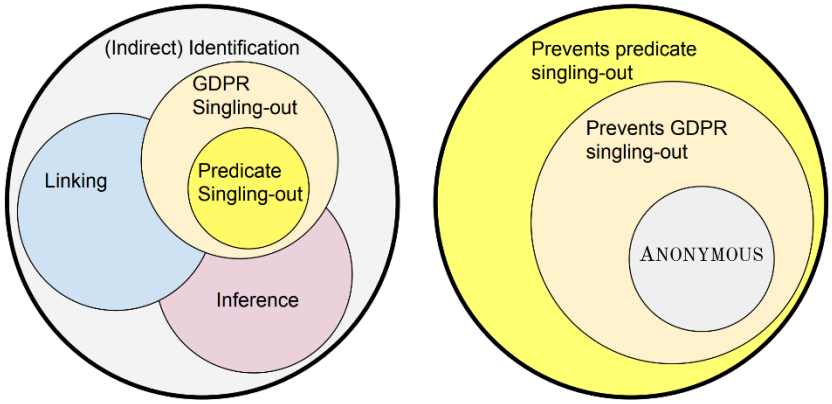

Figure 1: A graphical depiction of the proposed relationships between predicate singling-out and GDPR concepts (not to scale). On the left: linking, inference and (GDPR) singling out are three different ways of identifying personal data; predicate singling-out is a type of GDPR singling-out. On the right: for data to be considered ANONYMOUS, it must prevent GDPR singling-out and therefore also prevent predicate singling-out.

B. Modeling approach

A modeling of the GDPR notion of singling out should better apply to general anonymization mechanisms and attackers. That means that the modeling should make as few assumptions as possible regarding the operation of anonymization mechanisms and the strategies attackers may choose to use. In particular, the modeling should apply equally to existing anonymization mechanisms and to future ones, although the latter are not yet known, and, similarly, the modeling should equally apply to known attack strategies and to future ones, although the latter are also not yet known.

The modeling presented herein does not aim at perfectly representing the legal concept of singling out. Considering that singling out is just an example of "the means reasonably likely to be used" and therefore resilience to singling out is a necessary condition for ANONYMIZATION,[85] it would be most useful if the model would give a necessary condition for an anonymization mechanism to be considered secure against the GDPR notion of singling out (hence this condition would also be necessary for ANONYMIZATION).

Generality. The modeling presented here makes no assumptions about the format of an anonymization mechanism's output. Mechanisms are only assumed

[85] *See* discussion *supra* Section II.B.3.

to output *something* that an attacker might use to try to single out the input dataset.[86] For example, the output might be pseudonymized microdata, a quarterly earnings report, or an interface for querying a neural net. This is in contrast to approaches which only apply to very structured outputs, hence making a model narrowly applicable, as it would be inapplicable to other types of outputs which are prevalent in modern uses of data. An example of a model which is restricted to specific output formats is k-anonymity, which applies only to data in tabular form where generalization and suppression operations were applied.[87] As an example, an interface for querying a neural net cannot be said to be pseudonymous or k-anonymous; it is a category mistake.

The modeling herein also makes no assumptions about the attributes that an attacker can use to identify an individual within a dataset. In particular, an attacker is not assumed to use only so-called quasi-identifiers (also referred to as indirect identifiers), and may identify an individual using any combination of their attributes. For example, an attacker may identify an individual by specifying the collection of attribute values:

$$[\text{age} = 27, \text{gender} = \text{female, and income} \geq \$135{,}000].$$

The attacker may also specify a more involved property – a predicate – that the individual's attribute values satisfy, e.g.,

$$[\text{income} \geq \$5{,}000 \times \text{age}].$$

The modeling herein considers identification by predicates as falling under Article 4's notion of "indirectly" identifying natural persons.[88] The alternative to this modeling convention – namely, restricting attackers to identifying individuals by using only a collection of attribute values – would limit the scope of indirect identification in a way that ignores the complexities of modern data analysis. Looking ahead to Subsection III.C.3, the modeling will focus on predicates that express properties which are very rare in the underlying population.

Necessity. The modeling does not aim to perfectly represent the legal concept of singling out as a formal mathematical concept. The modeling attempts to

[86] Additionally, the modeling assumes that the anonymization mechanism is a fixed algorithm that takes as its only input a dataset to be protected. This setting directly captures general-purpose anonymization mechanisms such as mechanisms providing k-anonymity or differential privacy. If the anonymization mechanism requires additional inputs to anonymize successfully, then in many settings these can be thought of as being "hard-wired" or "fixed" as part of the mechanism, and hence these mechanisms also fit the modeling.

[87] *See* Ji-Won Byun, Ashish Kamra, Elisa Bertino & Ninghui Li, *Efficient k-Anonymization Using Clustering Techniques, 12* INT' L CONF. ON DATABASE SYS'S. FOR ADV. APPS. 188, 189 (2007) ("The k-anonymity model assumes that person-specific data are stored in a table (or a relation) of columns (or attributes) and rows (or records).").

[88] GDPR, *supra* note 6, at Article 4.

specify an important necessary condition for satisfying the requirements in the GDPR, rather than to perfectly model these requirements.[89] The resulting mathematical formalization hence lends itself to critiquing claims that certain privacy mechanisms offer the guarantees legally required in the GDPR. In other words, it can be used for disputing claims that such mechanisms are general-purpose ANONYMIZERS which transform personal data into ANONYMOUS data, and hence data which is outside the scope of the GDPR.[90]

The resulting model – a technical version of singling-out attacks called *predicate singling-out attacks* – implies the sort of singling-out attacks mentioned in Recital 26. Towards this end, predicate singling-out sets a high bar for an attacker to succeed. By making it difficult to qualify as a predicate singling-out attack, attention is focused on the most serious failures, whereas the GDPR notion of singling out may extend to other less serious failures. This bolsters the argument that any ANONYMIZATION mechanism must prevent predicate singling-out attacks.

The choice to prefer negative results is natural, at least in the context of evaluating GDPR ANONYMIZATION mechanisms. Because singling out is only one mode of identification under the GDPR, certifying that a mechanism prevents singling-out attacks under the GDPR would not suffice for certifying the mechanism as ANONYMIZING, as it ignores other failure modes. Conversely, if a mechanism permits singling-out attacks under the GDPR then it cannot ANONYMIZE.[91] As such, the legal implications of a model that can be used to demonstrate singling-out attacks are greater than a model that can be used to demonstrate the absence of singling-out attacks.

Finally, note that it is possible to prove that certain anonymization mechanisms do prevent predicate singling-out attacks.[92] While the significance of these claims is less clear, they do highlight some directions for future study.

C. Defining predicate singling-out

This Section defines predicate singling-out attacks and describes what it means for a privacy mechanism to prevent such attacks. It explains in detail the many modeling choices and other decisions that underlie the definition and provides intuition for the definition. Its structure parallels Section III.A, which provides a high-level summary. While this Section is essential for a deeper

[89] Regulations like the GDPR leave concepts such as singling out under-specified to allow for flexibility of interpretation in a wide range of situations, making perfect mathematical modeling an impossibility. *See* Cohen & Nissim, *supra* note 1, at 8345.

[90] A complementary approach is where the modeling attempts to specify a condition which is stringent enough that satisfying it suffices for satisfying a regulatory requirement. *See, e.g.*, Kobbi Nissim, Aaron Bembenek, Alexandra Wood, Mark Bun, Marco Gaboardi, Urs Gasser, David R. O'Brien, Thomas Steinke & Salil Vadhan, *Bridging the Gap between Computer Science and Legal Approaches to Privacy*, 31 HARV. J. L. & TECH. 689, 692 (2018).

[91] *See infra* Section IV.A.

[92] *See infra* Section IV.B.

understanding and appreciation of the methodology underlying this work, readers may wish to skim it during a first read.

1. The setting

Consider a setting in which a data controller has in its possession a *dataset D* consisting of personal information. In more detail, the dataset *D* is a table where each record contains the information of a single distinct individual.[93] Denote the number of records (i.e., distinct individuals) in the dataset *n* and the records of the dataset $x_1, x_2, \ldots, x_n$ (see Example 4 where $n = 3$).

The mathematical modeling refers to statements about information of individuals. Such statements, called *predicates*, reference variables corresponding to attributes that a dataset records about individuals. Examples of such variables include an individual's age, height, salary, and rent (see Example 4). A predicate expresses a condition on the possible attributes of an individual. It is a function which assigns a value True—meaning that a condition is satisfied—or False—meaning that a condition is not satisfied—to each possible individual. For convenience of reference, a predicate is often denoted by the letter *p*.

Example 4

A dataset D consists of $n = 3$ records:

	Name	age	height	salary	rent
x_1:	John Smith	49	6'0"	$70,000	$30,000
x_2:	Alice Liddell	32	5'6"	$100,000	$29,000
x_3:	Siobhan Murphy	50	5'5"	$50,000	$20,000

Consider the predicate p defined over the variables age, height, rent, and salary:

$$p: \text{[age = 49 AND height } \leq \text{5'8"] OR [rent } \leq \frac{\text{salary}}{3}\text{]}.$$

The predicate p assigns True if at least one of two conditions is met: either a person's age is 49 and their height is at most 5'8", or their rent is at most a third of their salary (or both). Applying the predicate to the dataset D, the predicate p assigns the value False to John's record: he is too tall and his rent is too high relative to his salary. Likewise on Siobhan's record. But p assigns True to Alice's records despite not satisfying the first clause of p due to her age: Alice's rent is less than a third of her salary. Because hers is the only record for which p is True, p isolates Alice in D.

[93] In real-world implementations, a dataset is often organized as a collection of several tables. Functionally, such datasets can in many cases be viewed as a single table with records corresponding to distinct individuals.

A predicate expresses a condition on a single record in the dataset without reference to the other records. So while they both uniquely describe John Smith, [height ≥ 6'0"] is a valid predicate but [the tallest person in D] is not.

Isolating using a data release. We consider an attacker attempting to single out using a data release generated by anonymizing a dataset, as described next (see Figure 2). To publish some version of the data in a dataset D, the data controller applies to D a computation that will be referred to as an *anonymization mechanism* and denoted M. The result of applying M to D is a *data release* $Y = M(D)$.[94] The data controller's goal is to publish a data release Y that is useful while preventing a potential *attacker*, who gets access to Y, from singling out. Note that it is the choice of the mechanism M which affects both the utility of the data release and the extent to which the data release makes singling out possible.

Once a data release Y is made public, it is available to potential attackers. For the purpose of this paper, an attacker applies to the data release Y an algorithm, denoted A. The attacker outputs a predicate p with the goal of singling out.

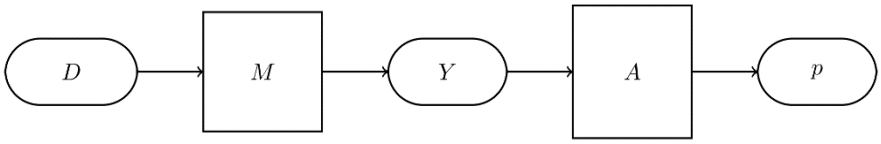

Figure 2: The anonymization mechanism M *is applied to the dataset* D *containing personal information to produce a data release* Y. *The data release can then be used by an attacker applying algorithm* A *in an attempt to produce a singling-out predicate* p.

The definition of singling out developed in the following Sections uses the notion of isolation as a starting point. A predicate p *isolates* a person in the dataset D if p evaluates to True on exactly one person's data x in D. For example,

[94] The formulation is very general: the anonymization mechanism M can be *any* algorithm that accepts personal data D as input and generates a purportedly ANONYMOUS output Y. For example, Y might be pseudonymized microdata, in which case M is the procedure for pseudonymizing the data in D. As another example, Y might instead be a neural network for facial recognition, with D a dataset of annotated headshots and M the machine learning procedures used to train the neural network. As such, this modeling applies with equal force to pseudonymized microdata, neural networks, aggregate statistics, regression models, or any other potential use of personal data. Note also that the mechanism M can be probabilistic, i.e., introduce randomness (colloquially, "noise") in its computation, e.g., for statistical disclosure limitation.

the predicate p in Example 4 isolates Alice Liddle in the dataset: it evaluates to True on x_2 and evaluates to False on all other records. An attacker *A isolates* a person in *D* using *Y* if it finds predicate *p* that isolates a person in *D*. In this case, the output of $A(Y)$ is a predicate *p* assigning the value True to exactly one record in *D* and assigning the value False to all the remaining records in *D*. Note that, as the attacker does not have direct access to the secret dataset *D*, it may be the case that the attacker does not immediately know whether or not the predicate *p* isolates in *D*, and hence whether it succeeded in isolating. However, it is not uncommon that the attacker would have means to verify whether the predicate *p* isolates in *D*, e.g., via access for a mechanism where it can query how many records in *D* satisfy *p*.

It is important to note that it is the data in the secret dataset *D* that needs to be protected, not the public data release *Y*. Hence the goal of the attacker is not to isolate a person in the public data release *Y* but to use the public data release to isolate a person in the secret dataset *D*. Furthermore, the alternative of isolating in the data release *Y* is incoherent when, as in the case of machine learning or highly aggregated statistics, there is no clear notion of an "individual's data" in *Y*. It is a category mistake to consider singling out on the release without making syntactic assumptions about it.

Modeling the data generation process and the attacker. The dataset *D* is modeled as consisting of n distinct records $x_1, \ldots, x_n$ each sampled independently from a single probability distribution. This is an accurate modeling for settings where data are collected via a well-designed small random sample of a large underlying population, such as the case when national statistical agencies run economic and social surveys. However, this modeling is a poor fit for many other settings which occur in practice, and, in particular, when data are collected about a tightly-knit community. Due to the similarities and the relationships between the community members, the data selected from distinct individuals cannot generally be considered independent from one another.

This modeling choice strengthens the argument that the modeling results in a definition of security against predicate singling-out attacks which is necessary for achieving ANONYMIZATION under the GDPR. First, if a mechanism *M* ANONYMIZES in all circumstances, then it should also ANONYMIZE in the special case that the input data are independently drawn from a distribution (for all possible distributions).[95] Second, an attacker's task is made harder under this assumption. For example, the modeled attacker cannot leverage information gleaned in prior data releases either by the data controller or by a third party, nor can any correlated data sources that may be publicly available be leveraged (e.g., correlated ratings on Netflix and IMDb). It cannot know about any individual

[95] That is, to demonstrate that a mechanism fails to ANONYMIZE, it suffices to exhibit a distribution and prove that the mechanism fails when the data is sampled according to the distribution. *See* Cohen & Nissim, *supra* note 1, at 8351.

records nor any correlations among records except insofar as they are revealed by *Y*.

On the other hand, the attacker (but not the data controller) is granted full knowledge of the data generation process. This makes the attacker's task easier, seemingly challenging necessity. However, attackers having some knowledge of the underlying distribution is a realistic assumption (indeed, this is the standard assumed in cryptography). Furthermore, we note that all the necessity claims made in this article only consider attackers which do not utilize any knowledge about the data generation process.[96]

Lastly, the modeling also grants the attacker knowledge of the mechanism deployed by the data controller, again seemingly challenging necessity. Notwithstanding, assuming a reasonably robust attacker is unavoidable because all mechanisms are protective against a sufficiently inept attacker. Further, in practice, many attackers have some knowledge of the mechanism and population distribution – roughly speaking, it is reasonable to assume attackers have general expertise but lack detailed insider knowledge of the specific contents of the data.

2. Isolation fails to capture singling out

Our goal is to define predicate singling-out so that for any anonymization mechanism *M* to be considered as preventing singling out under the GDPR, it must prevent predicate singling-out. As discussed in Section II.C, it is natural to define predicate singling-out as related to the risk that an attacker isolates a person in a dataset (see Example 2). Using the notation above, this would mean that the GDPR requirement of preventing singling out amounts to requiring that for every possible attacker *A* it would be very unlikely that the attacker *A* given the outcome of the anonymization mechanism $Y = M(D)$ successfully produces a predicate *p* which isolates a person in *D*.[97]

But this approach fails. If one makes no restriction on the types of predicates an attacker can use to isolate a person in a dataset, then no mechanism can make isolation improbable. Example 3 illustrates this phenomenon. In that example, the dataset *D* consists of 365 randomly selected students at a university. Because there are the same number of possible birthdays as people in the dataset, a randomly selected birthday has a decent chance of matching a single person. Indeed, the predicate [birthday = 03/12]—or any other date—has about a 37% chance

[96] Some of the attacks described in the following sections assume generated data contains a sufficient level of uncertainty, measured by the distribution's min-entropy. It is, however, not required for the singling-out attacker to know the distribution's min-entropy to mount the attack, or even to know that the distribution's min-entropy is sufficiently high (but the attack may fail in a case where the min-entropy is not sufficiently high). *See* Cohen & Nissim, *supra* note 1, at 8348.

[97] "Unlikely" here means that the probability that the predicate *p* successfully isolates a person should be very small. The probability is taken over all sources of randomness in this process, including the data generation process, and the randomness that the anonymization mechanism *M* and the attacker *A* may use in their computations. *See* Cohen & Nissim, *supra* note 1, at 8348.

of doing so. An attacker trying to isolate a person in this dataset could simply pick a random date and succeed with probability 37%. Example 5 explains that this is a general phenomenon: a 37% chance of isolation is always possible.[98]

A success probability of 37% is very high. If the first random date the attacker chooses fails to isolate a person, the attacker could simply try other dates, succeeding in short order. Ideally, the adversarial success probability should be very small, and should take into account that in today's computerized world, with automated attack attempts, an attacker may perform a huge number of attack attempts at a reasonably low cost. An attack success rate of, say, "just" 1% would mean that on average one of every hundred attack attempts would succeed, a rate that can amount to numerous successful attacks.

Returning to the example, such an "attack" should not be considered an attack at all. Note that the example does not even mention a mechanism M or data release Y. The same tactic succeeds against every M or none at all.

A privacy breach cannot be attributed to a data release that does not even exist. It should be a basic principle that the possibility of an attack should not be attributed to a particular use of data if it could be carried out just as well without it. The conclusion is that the notion of isolation alone fails to reasonably capture what GDPR singling-out may mean.

The task of identifying a predicate of weight $1/n$ was immediate with a dataset containing information of 365 students. Example 5 shows that this is a more general phenomenon. Mathematical tools exist for similarly identifying such predicates for data distributions that exhibit sufficient randomness, an amount of randomness which exists in typical data distributions.[99] Importantly, knowledge of the distribution is not needed for identifying a predicate of weight $1/n$.[100]

Example 5: Isolation by random guessing (generalization of Example 3)

Consider a dataset D consisting of n people's information, including sufficient detail so that each record is unique. Using *only* knowledge of the attributes included in the dataset and the population from which the data were drawn (and *nothing* else about D), it is straightforward to devise a predicate p that is `True` for about $1/n$ of the population. On average, one of the n people in the dataset will satisfy p. As in Example 3, a simple calculation shows that, with 37% chance, there is exactly one such person (see Technical Details 1). Hence, an attacker can successfully isolate a record in a dataset with 37% chance, even without seeing the dataset.

[98] *See* Cohen & Nissim, *supra* note 1, at 8346.

[99] *See id.*

[100] *See id.*

Technical Details 1: Calculation for Example 5

A calculation is included for the interested reader, but it can be skipped otherwise. For simplicity, the calculation neglects leap years and assumes that each birthday occurs with equal probability.

Consider first the chance that the predicate $[\mathsf{birthday} = 03/12]$ isolates the first person selected into the dataset. For that to happen, the first person selected into the dataset should have their birthday on $03/12$, an event that happens with probability $\frac{1}{365}$. Each of the other 364 persons selected into the dataset should have their birthday not on $03/12$, an event that happens with probability $1-\frac{1}{365}=\frac{364}{365}$. The probability that the predicate $[\mathsf{birthday} = 03/12]$ isolates the first person selected into the dataset is hence

$$\frac{1}{365}\cdot\underbrace{\frac{364}{365}\cdots\frac{364}{365}}_{364\text{ times}}=\frac{1}{365}\cdot\left(\frac{364}{365}\right)^{364}.$$

Similarly, the probability that the predicate $[\mathsf{birthday} = 03/12]$ isolates the second (or third, or fourth, …, or 365^{th}) person selected into the dataset is also $\frac{1}{365}\cdot\left(\frac{364}{365}\right)^{364}$.

Since these isolation events are mutually exclusive, as, by definition, a predicate cannot isolate more than one person, the probability that one of them occurs equals the sum of their probabilities, which amounts to the number of these events—365—times the probability of each event calculated above:

$$\underbrace{365}_{\#\text{ isolation events}}\cdot\underbrace{\frac{1}{365}\cdot\left(\frac{364}{365}\right)^{364}}_{\text{probability of each isolation event}}\approx 0.37.$$

More generally, a predicate which evaluates to True with probability $1/n$ isolates a person in a dataset with n records with probability

$$n\cdot\frac{1}{n}\cdot(1-\frac{1}{n})^{n-1}=(1-\frac{1}{n})^{n-1}>\frac{1}{e}\approx 0.37,$$

where $e\approx 2.71$ is Euler's number.

Finally, it is often the case that a predicate which evaluates to True with probability close to $1/n$ can be obtained based on knowledge of the

underlying distribution (similar to choosing an arbitrary date above to obtain a predicate which evaluates to 1 with probability $1/365$). However, a knowledge of the distribution is not generally needed to obtain such a predicate, as there exists a general technique for obtaining such predicates if the underlying distribution has a sufficient level of uncertainty (which is typically the case).[101]

3. Baseline risk of isolation

In contrast with the "attack" on isolation described in the previous Section, if an attacker uses the output of the anonymization mechanism $Y = M(D)$ to isolate a dataset in a way that would have been impossible without Y, then it is reasonable to "blame" the data release for making the isolation possible. Examples 6 and 7 illustrate such attacks.

Example 6: Perfect isolation

Consider a university dataset D consisting of the information of 365 students sampled randomly from a large state university, and a public data release Y created using an anonymization mechanism M. Suppose an attacker examines Y and—100% of the time—outputs a birthday that matches exactly one student. The attacker cannot be dismissed as merely getting lucky. The attacker *must* use the data release. Without it, the attacker can succeed at most 37% of the time.[102]

Example 7: Isolation with rare predicates

Consider the same student dataset D as in Example 6. This time, suppose the attacker succeeds in isolating a student with the predicate

p : [27-year-old] AND [Colombian] AND [Jewish] AND [vegan] AND [speaks Flemish] AND [pilot].

Individuals for whom this predicate evaluates to `True` (if they exist) are much rarer than individuals for whom the predicate [birthday = 03/12] evaluates to `True` (cf. Example 3). One would not expect to find even a single 27-year-old vegan Colombian Jewish pilot speaking Flemish in a population of many thousands or even millions, let alone a dataset of 365 students. Unlike the case of the birthday, an attacker can only isolate a

[101] *See also supra* note 96 and accompanying text.

[102] *See id.*

person using such a rare predicate—even if only once in a thousand times—if it knows $Y = M(D)$. This level of specificity can lead to a complete identification of the student by linkage to other datasets or via other methods (e.g., when combined with access to a mechanism answering statistical queries, as is discussed in Section III.E under Technical Details 2). It can also lead to a differential treatment of the student, even if their identity is not revealed.

More formally, whether an anonymization mechanism prevents predicate singling-out is related to whether an attacker can use the mechanism's output to isolate a person in a dataset more often than a statistical baseline. The statistical baseline is the highest probability of isolation any attacker can achieve without being provided with the output of the anonymization mechanism.[103] Unlike trivial isolation attacks that no mechanism can prevent, it is coherent to demand that an anonymization mechanism prevent isolation attacks that succeed much more often than the baseline.

Example 3 demonstrates that the statistical baseline can be as high as 37%. But the baseline is not a single number. As illustrated by Example 7, the baseline depends on how *rare* the predicate *p* is in the underlying population, where *p* is chosen by the attacker.[104]

One can precisely define a statistical baseline against which to measure an attacker's chance of successfully isolating a person in a dataset – a baseline that depends on how rare the attacker's predicate *p* is. The baseline is the probability that the best attacker isolates a dataset before seeing any data release. Generalizing the above examples, it is easy to isolate a record in a dataset *D* consisting of *n* people drawn from some population by using a predicate which matches about 1 in every *n* members of the population. The baseline probability of successful isolation remains about 37% whether the dataset is of size *n* = 365 or *n* = 365,000,000, as is analyzed in Technical Details 1. However, if the predicate matches drastically fewer than 1 in every *n* members of the population, the baseline risk becomes vanishingly small.[105]

A reasonable, and perhaps minimal, requirement from anonymization mechanisms which protect against singling out is that they would prevent attackers

[103] *See id.*

[104] *See id.*

[105] The exact relationship between the rarity of a predicate and the baseline risk can be mathematically quantified via a calculation similar to the one presented in Technical Details 1. For a predicate *p*, let *w* be the probability that a random person in the population satisfies *p*. The baseline level of risk is $\mathsf{baseline}(n, w) = n \cdot w \cdot (1-w)^{n-1}$. For example, $\mathrm{baseline}\left(1000, \frac{1}{1000}\right) \approx 37\%$, $\mathrm{baseline}\left(1000, \frac{1}{100{,}000}\right) < 1\%$, and $\mathrm{baseline}\left(1000, \frac{1}{10{,}000{,}000}\right) < .01\%$.

from isolating with rare predicates. This is the approach taken in the next Section in defining security against predicate singling-out.

4. Putting it all together – security against predicate singling-out

A *predicate singling-out attack* is defined to occur when an attacker who produces rare predicates succeeds in isolating a person in a dataset significantly more often than the corresponding baseline risk.[106]

An anonymization mechanism is defined to be *secure against predicate singling-out attacks* if it guarantees only an insignificant increase in the risk that an attacker can isolate a person when using very rare predicates.[107] Otherwise, we say the mechanism enables predicate singling-out attacks.

The definition as written requires guidance on setting thresholds to instantiate what "significantly more often" and "very rare" mean. First, what amount of risk increase is acceptable or unacceptable? Second, how rare must the attacker's predicates be? The exact tuning of the definition is ultimately a matter of policy, and we discuss some considerations next. However, even without setting these thresholds, the definition enables the quantitative analysis of and comparisons among the strength of various anonymization mechanisms against predicate singling-out attacks.

The more lenient the thresholds, the stronger the claim that the GDPR requires ANONYMIZATION mechanisms to prevent predicate singling-out attacks. By lenient, we mean that a predicate singling-out attack must be very effective to be considered a success.

Considering the risk increase over the baseline that is unacceptable, a strict policy would prohibit even a small increase in risk. A lenient policy would prohibit only very large increases (thus permitting some substantial increases). Recall that the baseline risk for very common predicates can be as high as 37%, while extremely rare predicates can have vanishingly small baseline risk.[108] Considering how rare the predicate must be, a strict policy would prohibit a large-enough increase in risk for rare and common predicates. A more lenient policy would only require the increase be small for exceedingly rare predicates.

Strengthening the argument that security against predicate singling-out is necessary for satisfying the requirements of the GDPR, we set both thresholds leniently. To successfully predicate single-out, an attacker must isolate a person using a vanishingly rare predicate. The predicate has to be rare enough that the chance that the predicate matches anyone in a different dataset drawn from the same distribution is essentially 0. For such rare predicates, the baseline risk is

[106] *See* Cohen & Nissim, supra note 1, at 8344.

[107] *See id.* at 8348.

[108] *See id*. at 8346.

also essentially 0. To successfully predicate single-out, we also require the attacker to succeed with significant probability – say, 10%.[109]

Note that restricting to such rare predicates does not make predicate singling-out impossible. People are unique. Any dataset describing people will admit a multitude of extremely rare predicates that isolate those people. At the same time, any fixed rare predicate will be very unlikely to match anybody in a new randomly-sampled dataset.

D. Example: A mechanism answering an exact count query

Datasets are often used to compute and release aggregate statistics such as counting how many people in the dataset meet some criterion. For example, how many people receiving unemployment benefits have been unemployed for less than 1 month? This Section works through Example 8, showing that a mechanism that releases a single exact count prevents predicate singling-out attacks.[110]

[109] Because the baseline risk is essentially 0, any significant level of isolation risk represents a significant increase over the baseline. The value 10% is chosen for illustration; all the results in this paper hold if we changed 10% to any number less than 36%.

[110] *See* Cohen & Nissim, *supra* note 1, at 8349.

Example 8: Advanced Physics prerequisites

Consider a database containing educational records of students in a high school. There are $n = 527$ students in the school, and hence the database contains 527 student records. Each student record includes identifying information such as name, ID number, date of birth, and home address; parent or guardian name and their contact information; courses taken, test scores, final grades, and academic status; medical and health records, including immunization records; disciplinary reports; documentation of attendance; previous schools attended; special education records; and more.

Planning for the next academic year, the school headmaster considers whether to open an Advanced Physics class for 12^{th} graders. The headmaster needs to know how many non-seniors meet the prerequisites for Advanced Physics: either Intro Physics or Intro Calculus (or both).

To do so, the headmaster needs to evaluate a *count query* on the database. A count query count_p is defined by a predicate p, and evaluating count_p on the database amounts to counting how many of the database records satisfy the predicate. In particular, the headmaster's predicate is

p : [NOT in 12th grade] AND
[passed Intro Physics OR passed Intro Calculus].

Applying the predicate p to each student record in the database results in True if the student is not in 12^{th} grade and has passed Intro Physics or Intro Calculus (or both classes). Otherwise, applying p to the record results in False.

Evaluating count_p on the database would result in the count C: the number of database records for which p evaluates to True (i.e., the number of students eligible to enroll in Advanced Physics.)

Count queries are an essential ingredient of data analysis, statistics, and machine learning, and they receive extensive treatment in the mathematical literature on private data release.[111] Do count queries enable or prevent predicate singling-out attacks?

[111] Many papers describe count queries and differential privacy in the context of count queries. For examples of some of the early work in this area, see, e.g., Stanley L. Warner, *Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias*, 60 J. AM. STATISTICAL ASS'N 63 (1965); Avrim Blum, Cynthia Dwork, Frank McSherry & Kobbi Nissim, *Practical Privacy: The SuLQ Framework*, 24 PROC. ACM SIGMOD-SIGACT-SIGART

Example 8 (cont'd)

If she learns the count C, the headmaster will have to include it in a publicly available report to the school board. She tries to reason whether releasing this count would enable a predicate singling-out attack against the student database.

Learning the count C could certainly help an attacker *isolate* a student in the database. By learning some information about how many students satisfy the predicate p used in the count query count_p, the attacker gains some additional information useful for isolating a student in the school database. For example, if there were only $C = 2$ eligible students, an attacker who wished to isolate a single student might guess that there is only a single eligible *male* student. This attacker would try to isolate using p_M:

p_M : [NOT in 12th grade] AND
[passed Intro Physics OR passed Intro Calculus]
AND [Male].

But predicate singling-out requires more than just isolating a student. It demands that the predicate that isolates must be vanishingly *rare* – so rare that one would be extremely surprised to learn that *any* students at the school satisfied it. The predicate p_M is not nearly so rare. It matches all male high school seniors eligible to take Advanced Physics. (Indeed, if it was rare enough to count towards predicate singling-out, the headmaster would have little reason to make the original count query in the first place: she would already know the answer is almost certainly 0.)

The example illustrates a crucial difference between predicate singling-out and mere isolation: the rarity of the predicates considered. Note that the rarity of the predicate p_M in the example is not determined by how many students in the school satisfy it. The rarity is determined by how many students in the school

SYMP. ON PRINCIPLES DATABASE SYS. 128 (2005); Cynthia Dwork, Frank McSherry, Kobbi Nissim & Adam D. Smith, *Calibrating Noise to Sensitivity in Private Data Analysis*, 3 THEORY CRYPTOGRAPHY CONF. 265 (2006); Cynthia Dwork, Krishnaram Kenthapadi, Frank McSherry, Ilya Mironov & Moni Naor, *Our Data, Ourselves: Privacy via Distributed Noise Generation*, 25TH INT'L CRYPTOLOGY CONF. (EUROCRYPT 2006) 486 (2006).

are *expected to satisfy it a priori*.[112] The predicate is rare if the prior expected number is 0 or extremely close to it.

The specific attacker above failed to leverage the count query into a predicate singling-out attack. As shown next, this is inherent: a single count prevents predicate singling-out attacks in general.

The argument is as follows: to predicate single-out, an attacker A must – with probability at least 10% – find a predicate p^* that is very rare but isolates a person in the school database.[113] To show that the exact count C does not enable predicate singling-out attacks, it suffices to prove the following: for any attacker A that takes as input C and predicate singles-out, there is a *trivial* attacker B which predicate singles-out nearly as often as A *without using* C. This suffices because of two observations. First, the baseline risk for very rare predicates is negligible – much smaller than one in a million. Second, the baseline risk is defined as the maximum success probability of *any* trivial attacker (those that get no information about the database), including B. Putting it all together, A does not succeed much more often than B, which itself does not do any better than the baseline, which in turn is much, much smaller than one in a million for very rare predicates. Therefore A does not succeed much more often than once in a million, a far cry from the necessary 10%.

Example 8 (cont'd)

Following the above argument, the headmaster must show that for any attacker A that uses the count C to predicate single-out, there is a trivial attacker B that predicate singles-out nearly as often without using C.

She first observes that, as there are 527 students, the answer C to the count query must be one of the 528 values $0,1,2,\ldots,527$. Given C, the attacker A would output a predicate p^* which it hopes is both extremely rare and isolates an individual student in the school.

She then considers the hypothetical trivial attacker B (that does not get to know C or anything else about the database). The trivial attacker B would act as follows. First, it will guess at random an "answer" to the query, i.e., a random guess g in the range $0,\ldots,527$. Then it will act exactly as the attacker A would with the exception that, where A expects the true count C as an answer to the count query, B would use its guess g instead. In other words, B would replicate A's attack on the dataset, except that B would use the random guess g instead of querying the dataset for the count C.

[112] Formally, this only makes sense if there is some underlying distribution from which the students are drawn. *See supra* Subsection III.C.1. One of the limitations of the predicate singling-out framework is that, as in the example, this assumption often does not hold.

[113] *See* Cohen & Nissim, *supra* note 1, at 8349.

To continue with her analysis, the headmaster notes that, as there are 528 values in the range $0, \ldots, 527$, the guess g happens to equal C with probability $1/528$. Furthermore, in the event where the guess g happens to equal C, the trivial attacker B does exactly what A does when given the real answer C. Hence, the trivial attacker B succeeds with probability which is *at least* $1/528$ times the probability A succeeds.

Equivalently, A succeeds in its attack with probability which is *at most* 528 times the probability the trivial attacker B does. Because B's probability of success is smaller than the baseline which is itself negligibly small, A's probability of success is also negligibly small.

The headmaster can publish the exact number of students eligible to take Advanced Physics without enabling attackers to predicate single-out against the school's database.

The argument outlined above holds in general. Thus, a mechanism that exactly answers a single count query prevents predicate singling-out attacks.[114]

E. Composability

One of the benefits of having a rigorously defined notion is that that notion can itself be examined rigorously, so as to understand its strengths and weaknesses. As we describe next, security against predicate singling-out attacks lacks *composability*, an important goal of theoretically sound mathematical privacy concepts. This strongly suggests that preventing these attacks does not by itself guarantee individual privacy nor achieve full ANONYMIZATION.

All data uses degrade data privacy.[115] It is impossible to make unlimited use of data without risking its disclosure. The best one can hope for is that the extent of the degradation – even across many uses – can be quantified and thereby limited.

Composability makes it possible to quantify the degradation resulting from many uses by quantifying and combining the degradation from each use individually.[116] Privacy concepts like predicate singling-out, *k*-anonymity,[117] and differential privacy[118] provide competing ways of quantifying privacy degradation. Some privacy concepts, like differential privacy, are composable: the privacy loss due to many independent differentially private data analyses on a

[114] For the formal analysis (including a formal definition of negligible quantities), see Cohen & Nissim, *supra* note 1, at 8347.

[115] *See* Fluitt et al., *supra* note 5.

[116] *See id.* at 3.

[117] *See infra* Section IV.A and references therein.

[118] *See infra* Section IV.B and references therein.

single data set can be quantified.[119] Others, like *k*-anonymity, are not: the combination of just two *k*-anonymous data analyses may fail to provide any level of protection.[120]

Security against predicate singling-out attacks is not composable. Like *k*-anonymity, the combination of just two such secure mechanisms may enable predicate singling-out attacks.[121] As such, whether a mechanism (only) prevents predicate singling-out attacks should not be conflated with whether the mechanism is protective of individual privacy. We view it only as a necessary condition for ANONYMIZATION, not a sufficient condition. To the extent that predicate singling-out attacks capture the essence of GDPR singling-out, the legal concept may also lack composability.

Technical Details 2

A sketch of the argument against composability is included for the interested reader, and can be skipped otherwise.[122] We describe how a collection of exact count queries can be used to predicate single-out. (Exhibiting a failure of composability with just two mechanisms is more complicated.)

We build on Example 8, allowing the headmaster to make arbitrary exact count queries and publish her results. For any predicate p, she can learn count_p: the number of students satisfying p. The headmaster's strategy is to use count queries to learn a highly detailed description of a single student. From there, predicate singling-out is easy: the detailed description is itself a sufficiently rare predicate that isolates with probably 100%.

The headmaster notes that, as in Example 5, it is easy to come up with a predicate p_{isolate} that would isolate in the school database with probability close to 37%. A successful isolation occurs when the answer 1 is returned for the query $\mathsf{count}_{p_{\mathsf{isolate}}}$. Because the headmaster can check whether or not a given predicate isolates, she can easily find one that does.

Knowing a predicate p_{isolate} that isolates, it is easy to modify it to gain more information on the isolated student. For example:

$$p_{absent} : p_{isolate} \text{ AND } [\geq 15 \text{ absences}].$$

[119] *See* Wood et al., *supra* note 28, at 244.

[120] *See* Srivatsava Ranjit Ganta, Shiva Prasad Kasiviswanathan & Adam Smith, *Composition Attacks and Auxiliary Information in Data Privacy*, ACM SIGKDD INT'L CONF. ON KNOWLEDGE DISCOVERY AND DATA MINING 265 (2018).

[121] *See* Cohen & Nissim, *supra* note 1, at 8349.

[122] *See* Cohen & Nissim, *supra* note 1, at 8349.

The answer to the query $\text{count}_{p_{\text{absent}}}$ is 1 if the isolated student has at least 15 absences and 0 otherwise. In this way, the headmaster can learn virtually any information that is stored in the education records of the student isolated by p_{isolate}. As explained above, this allows her to predicate single-out the student.

In conclusion, the headmaster realizes that there exist collections of count queries which result in predicate singling-out. She may inadvertently make such queries, or may be influenced to make such queries by someone who wishes to single-out in the school database.

IV. ANALYZING MECHANISMS USING PREDICATE SINGLING-OUT

This Section demonstrates how the definition of security against predicate singling-out can help analyze specific privacy technologies and make informed decisions about whether the use of these privacy technologies satisfies the GDPR's ANONYMIZATION requirement.

A. k-anonymity

k-anonymity is a popular approach to anonymizing microdata introduced by Sweeney and Samarati.[123] *k*-anonymity is 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."[124]

The Working Party opinion on anonymization techniques states that "*k*-anonymity techniques aim to prevent a data subject from being singled out by grouping them with, at least, *k* other individuals."[125] The opinion then concludes that *k*-anonymity prevents GDPR singling-out attacks.[126]

This Section provides two examples of predicate singling-out attacks that are possible against *k*-anonymous mechanisms. Because we consider preventing predicate singling-out attacks to be a necessary precondition for preventing GDPR singling-out attacks, our conclusions challenge this opinion of the Working Party.

Notably, and unlike most previous attacks against *k*-anonymity, predicate singling-out attacks can be carried out even if *every* attribute is classified as a quasi-

[123] *See* Pierangela Samarati & Latanya Sweeney, *Protecting Privacy when Disclosing Information: k-Anonymity and Its Enforcement through Generalization and Suppression*, Working Paper (1997).

[124] Sweeney, *supra* note 15, at 557.

[125] *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 16.

[126] *See id.* at 24.

identifier subject to the *k*-anonymity guarantee. As such, they apply equally to *ℓ*-diversity and *t*-closeness, variants of *k*-anonymity also deemed effective against singling-out by the Working Party.[127]

The example attacks illustrate that as a justification of ANONYMIZATION, *k*-anonymity falls short. This is not to say that *k*-anonymous data releases are always vulnerable to attack. When created by experts with sufficient care, *k*-anonymous data releases may not allow singling-out. However, the protection is perhaps better attributed to the practitioner's expertise and care.[128]

1. *k*-anonymity background

A dataset is *k*-anonymized by applying suppression and generalization operations to certain attributes called *quasi-identifiers*.[129] Quasi-identifiers are those attributes that may be available from other sources and thereby may be used to identify individuals, even if indirectly.[130] *k*-anonymity requires that what is revealed about an individual's quasi-identifiers should be the same as what is revealed about at least $k - 1$ other individuals' *quasi-identifiers*.[131] Other attributes – including highly sensitive attributes – are left intact. The number *k* is a parameter of the definition, with higher *k* thought to offer greater privacy. Typical choices include $k = 5$ and $k = 10$.

[127] *See id.* at 18.

[128] For practitioners we offer only preliminary advice: global recoding using a pre-defined, data-independent hierarchy may be less vulnerable to predicate singling-out attacks.

[129] *See* Samarati & Sweeney, *supra* note 123, at 1.

[130] *See* Sweeney, *supra* note 15, at 563 ("[I]t is usually publicly available data on which linking is to be prohibited and so attributes which appear in private data and also appear in public data are candidates for linking; therefore, these attributes constitute the quasi-identifier and the disclosure of these attributes must be controlled.").

[131] Samarati & Sweeney, *supra* note 123, at 4.

Example 9

The following is an example of a dataset with six records and a k-anonymized version of the same dataset for $k = 2$. In this example, ZIP code, age, and sex are considered to be quasi-identifiers. The sensitive attribute disease is not considered a quasi-identifier and is kept unchanged. An asterisk denotes a symbol which has been suppressed.

D:

ZIP	Age	Sex	Disease
23456	55	F	COVID
23456	42	F	COVID
12345	30	F	Emphys.
12346	33	M	Lung
13144	45	F	Heart
13155	42	M	Hepatitis

Y:

ZIP	Age	Sex	Disease
23456	**	*	COVID
23456	**	*	COVID
1234*	3*	F	Emphys.
1234*	3*	F	Lung
131**	4*	*	Heart
131**	4*	*	Hepatitis

Each record in Y is derived from a record in D. Y is also 2-anonymous: each record's quasi-identifiers are the same as one other record's.

Suppose an attacker knows that Jeanine – a female living in 13144 – is included in the dataset. Given Y, the attacker who knows nothing else cannot determine which of the last two records corresponds to Jeanine. On the other hand, if Jeanine lived in 23456, the attacker can be sure that she suffered from COVID. Moreover, if Jeanine lived in 12345, the attacker can conclude that she is (or was) probably a smoker. These are examples of "homogeneity attacks."

It is well established that k-anonymity suffers from a number of weaknesses.[132] As in Example 9, a "homogeneity attack" can reveal sensitive information about a target person known to be in the dataset, even if an attacker cannot determine which record in Y corresponds to the target.[133] This and other attacks have motivated a slew of variant definitions, including ℓ-diversity[134] and t-closeness.[135]

2. Example predicate singling-out attacks

We provide two examples of predicate singling-out attacks on k-anonymous datasets. The first, Example 10, is rather contrived – i.e., experts applying k-anonymity would be very unlikely to use a mechanism like the one described in that example. It is presented here only for the simplicity of its illustration. The second, Example 11, is a more complex argument but applies for typical k-anonymous data releases when there are many quasi-identifiers.[136] It illustrates how a k-anonymous mechanism can assist the predicate singling-out attacker by providing an almost-isolating low-weight predicate; the predicate is easily modified into a predicate which is both isolating and low-weight, as required to predicate single-out.

Example 10: Predicate singling-out from data-dependent intervals

Note that, although experts applying k-anonymity would be very unlikely to use the k-anonymous mechanism in this example, this example is useful as a simplified illustration of a singling-out attack on a k-anonymous dataset.

[132] *See generally* Ashwin Machanavajjhala, Daniel Kifer, Johannes Gehrke & Muthuramakrishnan Venkitasubramaniam, *ℓ-diversity: Privacy beyond k-anonymity*, 1 ACM TRANSACTIONS ON KNOWLEDGE DISCOVERY FROM DATA 3 (2007) (describing a number of different types of attacks on k-anonymity).

[133] *See id.* at 3-4.

[134] *See id.* at 3.

[135] *See generally* Ninghui Li Tiancheng Li & Suresh Venkatasubramanian, *Closeness: A New Privacy Measure for Data Publishing*, 22 IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENG'G 943 (2010).

[136] By typical, we roughly mean that the published quasi-identifier values that define the equivalence classes in the k-anonymous release themselves define an exceedingly rare predicate (meeting the rarity threshold in the definition of predicate singling-out). *See* Cohen & Nissim, *supra* note 1. Many k-anonymous mechanisms (especially those using "local recoding") have this property when there are many quasi-identifiers. *See* Aloni Cohen, *The Quasi-identifiers are the Problem: Attacking and Reidnetifying k-Anonymous Datasets* (forthcoming). In fact, local-recoding mechanisms enable much stronger predicate singling-out attacks that use very rare predicates and succeed in isolating 99% of the time. *See id.*

Consider a dataset consisting of $n = 50$ American passport holders and their passport numbers $x_1, x_2, \ldots, x_{50}$ in ascending order. For our stylized example, assume that passport numbers are randomly-selected 9-digit numbers. The data controller considers a passport number to be a quasi-identifier, and hence k-anonymizes the dataset as in the table below with $k = 10$. The 10-anonymized dataset consists of a single column which contains ranges of passport numbers:

	Passport number
records 1 to 10:	between x_1 and x_{10}
records 11 to 20:	between x_{11} and x_{20}
records 21 to 30:	between x_{21} and x_{30}
records 31 to 40:	between x_{31} and x_{40}
records 41 to 50:	between x_{41} and x_{50}

Publishing the above table directly exposes the passport numbers of ten people in the dataset, namely x_1, x_{10}, x_{11}, x_{20}, x_{21}, x_{30}, x_{31}, x_{40}, x_{41} and x_{50}. With this observation, singling out is straightforward. The attacker simply outputs the predicate

$$p\text{:} \quad [\text{passport number} = x_1].$$

To predicate single-out, p must be extremely rare and isolate with probability at least 10%. For this example, we assumed that passport numbers are uniformly at random from all one billion possible 9-digit numbers. The predicate is sufficiently rare: the chance that a new passport number satisfies the predicate p (i.e., is equal to x_1) is one in a billion ($\frac{1}{10^9}$). The probability of isolation is very close to 100%. (If p does not isolate, then one of x_2, …, x_{50} satisfies p; this occurs with probability at most $\frac{49}{1{,}000{,}000{,}000}$, which is less than one in ten million.)

The simple mechanism in Example 10 enables predicate singling-out: it unambiguously reveals some of the underlying data. Contrast this with the ZIP codes in Example 9, which were generalized in a way that leaves some uncertainty as to the true underlying values. The fact remains that *k*-anonymity does not necessarily prevent predicate singling-out attacks and may even enable them. Example 11 reinforces this conclusion.

Example 11

Consider a dataset where the quasi-identifiers include a list of dates of doctor visits. The equivalence classes in the table are k-anonymized using date ranges, resulting in the following table:

	Date 1	**Date 2**	**. . .**	**Date 30**
records 1 to k:	Jan 1, 2017	Feb 1–10, 2017	. . .	Nov–Dec, 2019
$k+1$ to $2k$:	Dec 5, 2016	Jan–Feb, 2017	. . .	2019–2020
. . .	. . .	. . .	. . .	. . .
last k records:	Mar 15–31, 2018	Apr 6, 2018	. . .	Dec 25, 2018

Observe that this table already contains very specific information about people in the dataset. Namely, we know that k people satisfy the predicate

$$p\text{: [Date 1} = \text{Jan 1, 2017] AND [Date 2 } \textit{in}\text{ Feb 1 – 10, 2017] AND ... AND [Date 30 } \textit{in}\text{ Nov – Dec, 2019].}$$

Depending on the data distribution, this predicate will be exceedingly rare – rare enough to meet the threshold for a predicate singling-out attack. (To be sure, at least k people in the dataset satisfy it. But a new random person drawn from the distribution is very unlikely to satisfy it.)

Therefore, any more specific predicate will also be exceedingly rare. All the attacker must do is find a more specific predicate that has a good chance of isolating somebody in the dataset. For example, the attacker can guess that Date 2 is in fact Feb 1, 2017. The resulting predicate is

$$p'\text{: [Date 1} = \text{Jan 1, 2017] AND [Date 2} = \text{Feb 1, 2017] AND ... AND [Date 30 } \textit{in}\text{ Nov – Dec, 2019].}$$

If $k = 10$, then under reasonable assumptions on the data distribution, the probability that this new predicate isolates a single person is about 37%. Using p', the attacker succeeds in isolating more than 10% of the time, meeting the other threshold for a predicate singling-out attack.

It is no coincidence that 37% appears in Example 11. From the published data, an attacker can easily infer a rare predicate p that matches only k data subjects.

All that remains is to isolate a single person from this group of k using any (even very common) predicate. Random guessing succeeds with probability 37%.[137]

B. Differential privacy

Differential privacy is a definition (or standard) of privacy introduced by Dwork, McSherry, Nissim, and Smith in 2006.[138] The definition applies to analyses performed over collections of individual data. It articulates in precise mathematical language a requirement to protect the information pertaining to each individual in the data: the outcome of the analysis should be randomized (i.e., it can provide different results if executed twice on the same data),[139] and its outcome distribution should not change significantly whether the individual's information is or is not included in the input to the analysis.[140]

This Section explains that differential privacy prevents predicate singling-out attacks. Moreover, differential privacy prevents a much wider class of attacks than strictly required by our definition.[141] And while security against predicate singling-out attacks does not necessarily compose,[142] the protection afforded by differential privacy does compose. These findings are consistent with the Working Party opinion on anonymization techniques.[143]

As described in Section III.C, predicate singling-out is more useful as a tool for critiquing anonymization mechanisms than for defending them. Whereas we view a mechanism's enabling of predicate singling-out attacks as strong evidence against ANONYMIZATION, it is less clear how to interpret the finding that differential privacy prevents attacks. Our best guess is that differential privacy probably prevents GDPR singling-out, not just predicate singling-out. In any case, the result suggests that differential privacy shows some promise as an

[137] *See* Cohen & Nissim, *supra* note 1, at 8352; *see also supra* Example 5.

[138] *See* Cynthia Dwork, Frank McSherry, Kobbi Nissim & Adam Smith, *Calibrating Noise to Sensitivity in Private Data Analysis*, PROC. OF THE THIRD THEORY OF CRYPTOGRAPHY CONFERENCE 265 (2006). For a non-technical introduction to differential privacy and a reading list on the topic, see generally Wood et al., *supra* note 28.

[139] *See* Wood et al., *supra* note 28, at 233.

[140] See Wood et al., *supra* note 28, at 212. Throughout this section on differential privacy, the unit of protection is an individual's entire record (sometimes referred to as *user level* differential privacy). The argument that differential privacy entails security against predicate singling-out (presented below) does not hold when the unit of protection is an individual's (single) event or action. In particular, event or action level differential privacy does not guarantee security against predicate singling-out. This means that some event or action level differentially private mechanisms can be demonstrated to fail to be secure against predicate singling-out.

[141] This corresponds to different settings of the thresholds in Subsection III.C.4.

[142] *See supra* Section III.E.

[143] *See Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 15 ("If only statistics are output and the rules applied to the [dataset] are well chosen, it should not be possible to use the answers [of a differentially private mechanism] to single out an individual.").

ANONYMIZATION technique. We emphasize that preventing singling out is not by itself enough for ANONYMIZATION, and we do not mean to suggest that differential privacy achieves ANONYMIZATION.

1. Differential privacy background

A full definition of differential privacy is beyond the scope of this article. We briefly describe a few of differential privacy's most salient and relevant features.[144]

First, differential privacy is parameterized.[145] The definition of differential privacy is equipped with a parameter, often denoted by the Greek letter ε.[146] The parameter ε is set to a positive number and puts a limit on privacy loss.[147] A smaller ε (e.g., 0.1) corresponds to a stricter privacy requirement, and also to stronger limitations on the utility of the computation, measured, e.g., by the accuracy or number of computations that can be performed.[148] Intuitively, by setting ε, one directly expresses a limit on the total privacy loss allowed, and indirectly imposes limits on the computation accuracy.[149]

Second, differentially private computations are noisy.[150] To mask the differences between the outcome of an analysis with and without an individual's information being included, a differentially private analysis must introduce some amount of randomness. A choice of a smaller ε corresponds to a better masking, which in turn implies larger noise.[151] Statistical analyses, performed with differential privacy hence differ from standard statistical analyses.

Third, differential privacy composes.[152] As discussed in Section III.E, composability, i.e., that a combination of mechanisms satisfying a security requirement should also satisfy the requirement (usually with weaker parameters), is a desirable property for any security definition. For differential privacy, a mechanism combined from several differentially private mechanisms is itself differentially private (albeit, with weaker parameters).

Example 12: Advanced Physics prerequisites with differential privacy

A differentially private mechanism for answering count queries must introduce statistical noise to its answer, returning an approximate count

[144] For an in-depth discussion of the definition of differential privacy and its semantics, see Wood et al., *supra* note 28.

[145] *See id.* at 234.

[146] *Id*. at 212.

[147] *Id.*

[148] *See id*. at 235-36.

[149] *See id.*

[150] *See id.* at 233.

[151] *See id*. at 235-36.

[152] *See id.* at 244.

instead of the exact count. For instance, using what is known as the "Laplace Mechanism," the approximate count is computed by adding a specific type of random error to the true count. If ε is set to 0.5, the error is at most ± 5 with 95% probability.

Recall Example 8, where a headmaster wishes to use a count query count_p to determine how many seniors meet the prerequisites for Advanced Physics. Suppose that there are 28 students in the high school who meet the prerequisites. The school headmaster evaluates the query count_p using the differentially private Laplace Mechanism and receives the answer "31.2". Taking into account the mechanism and the choice of $\varepsilon = 0.5$, the headmaster concludes that the number of students eligible to enroll in the Advanced Physics class is close to 31 and very likely between 26 and 36 – an estimate which, while not exact, would allow her to continue with the planning.

If the school headmaster makes another differentially private count query – say, to estimate the number of students with a cumulative GPA of 3.5 or higher – the privacy guarantees degrade. It must: more information is simply more disclosive. However, because both count queries were made using differential privacy, it is possible to quantify the degradation. For example, if both count queries made using $\varepsilon = 0.5$, the "effective ε" of the combined counts is at most 1.

2. Differential privacy prevents singling-out attacks

Differential privacy prevents predicate singling-out attacks. The mathematical proof of this fact is somewhat involved.[153] Because differential privacy itself composes, its protection against predicate singling-out attacks also composes. Furthermore, when using differential privacy, the resulting security against singling out will compose because differential privacy itself composes. This

[153] *See* Cohen & Nissim, *supra* note 1, at 8351. The proof relies on the recently discovered connection between differential privacy and statistical generalization. *See, e.g.*, Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold & Aaron Roth, *Preserving Statistical Validity in Adaptive Data Analysis*, PROC. FORTY-SEVENTH ANN. ACM SYMP. ON THEORY COMPUTING 117 (2015); Raef Bassily, Kobbi Nissim, Adam Smith, Thomas Steinke, Uri Stemmer & Jonathan Ullman, *Algorithmic Stability for Adaptive Data Analysis*, PROC. FORTY-EIGTH ANN. ACM SYMP. ON THEORY COMPUTING 1046 (2016). Very roughly, any correlations (e.g., smoking and emphysema) discovered using a dataset sampled from a large population will generalize – apply almost as well in the population as a whole – if the data analysis is differentially private. An implication is that an attacker cannot use the outcome of a differentially private analysis to come up with a very rare predicate that matches *anybody* in the dataset (as required to predicate single-out).

provides one way to circumvent the difficulties presented by the non-composition of security against singling out.

However, it is easy to see that there are mechanisms that are not differentially private but do prevent predicate singling-out attacks. One exact count query prevents predicate singling-out attacks.[154] Differential privacy requires noise; it cannot be exact.[155] As such, any mechanism that outputs an exact count is not differentially private.[156]

The protection offered by differential privacy goes further than our definition requires. Namely, differential privacy prevents predicate singling-out attacks where the attacker's predicate only has to be slightly rare, not very rare. It also provides a precise quantitative bound on the increase over the baseline risk as a function of the predicate's rarity and the differential privacy parameter ε.

Differential privacy also highlights the danger of conflating the goal of preventing predicate singling-out attacks (with extremely rare predicates) and that of achieving ANONYMIZATION. The definition is intended as a necessary condition for ANONYMIZATION, not a sufficient one. Differential privacy offers evidence against sufficiency. Namely, it satisfies the definition even when the privacy parameter is very large (e.g., $\varepsilon = 50$), a regime where differential privacy provides scant protections. The significance of the choice of privacy parameter was also recognized by the Working Party, further suggesting that ANONYMIZATION requires smaller ε.[157] This issue can be solved by considering different choices of the thresholds discussed in Subsection III.C.4.

V. WHAT PREDICATE SINGLING-OUT MEANS FOR GDPR SINGLING-OUT

This Section closes a cycle between predicate singling-out, GDPR singling-out, and ANONYMIZATION. Parts II and III began with the legal concepts of ANONYMIZATION and GDPR singling-out and derived the mathematical concept of predicate singling-out. This Section returns from the technical concept to the legal concept and enumerates insights gained from predicate singling-out and how they apply to GDPR singling-out and ANONYMIZATION. It also considers different ways of resolving the tension between the findings described herein and the guidance provided by the Working Party.

---

[154] *See supra* Section III.D.

[155] For a discussion explaining why differential privacy requires noise addition, see Wood et al., *supra* note 28, at 232–37.

[156] Furthermore, the security against predicate singling-out attacks provided by exact counts does not compose, *see supra* Section III.E, whereas the security provided by differential privacy does, *see* Wood et al., *supra* note 28, at 244.

[157] *See Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 16 (identifying "not injecting enough noise," which corresponds to using large ε, as a "common mistake").

A. Insights from the mathematical formulation

Policymakers wishing to regulate singling-out attacks should regulate predicate singling-out attacks. If, as this article argues, predicate singling-out does succeed in capturing a class of GDPR singling-out attacks, the technical concept provides a powerful new test for evaluating ANONYMIZATION techniques. Any mechanism that purports to ANONYMIZE personal data in all circumstances must necessarily prevent predicate singling-out attacks. Data protection authorities can evaluate the offerings of ANONYMIZATION-as-a-service companies[158] and provide an unambiguous minimum threshold test that any ANONYMIZATION product must clear. At a minimum, regulation should cover the most blatant attacks: those with a very large increase over the baseline risk using very rare predicates (Part III). Future guidance on anonymization techniques would also have to revise the Working Party's finding that *k*-anonymity generally prevents GDPR singling-out. A more expressive policy might assign to every rarity an acceptable maximum risk increase.

Attempting to precisely define a type of GDPR singling-out attack gives new insight into the concept, and there are lessons to be learned from the analysis in Parts II–IV about GDPR singling-out (and hence ANONYMIZATION) even if predicate singling-out does not exactly capture its meaning.[159]

First, one should be careful with respect to correctly identifying the data which needs to be protected against singling out. With the deployment of paradigms such as *k*-anonymity, where data is suppressed or generalized but is otherwise left in its original form, it can be tempting to assume that it is the published *k*-anonymized data which needs to be protected against singling out.[160] Rather, it is the original dataset to which anonymization was applied that needs to be protected, and hence whether singling out occurs needs to be checked with respect to it. Perhaps a misunderstanding on this point formed the basis for the Working Party's statement that *k*-anonymity prevents singling out.

Second, baseline risk, probability of attack, and the rarity of a predicate are all relevant to singling out. Baseline risk – i.e., how well an attacker can do without access to the data – is fundamental. It should be a basic principle that the possibility of an attack should not be attributed to a particular use of data if it could be carried out just as well without it. This principle applies not only to

[158] *See, e.g.,* AIRCLOAK, https://aircloak.com [https://perma.cc/YZ4X-WMSG]; PRIVITAR, https://www.privitar.com [https://perma.cc/RY4Y-MJA8]; ANONOS, https://www.anonos.com [https://perma.cc/93CD-B3LD].

[159] For example, consider the possibility of a future European Data Protection Board decision or guidance document that contradicts the interpretation of singling out that appears in *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 11.

[160] With such a view, it would indeed be impossible to single out in a *k*-anonymized data, but also in the outcome of an anonymizer which does not suppress or modify its input, but simply duplicates every individual's record.

singling out, but also to linkability and inference.[161] Because of this principle, the idea that singling out is the same as isolation – an idea that may seem intuitive and appealing – had to be rejected (Subsection III.C.2); it led to considering the rarity of a predicate and to defining the susceptibility to attacks in terms of probability. What makes an attack meaningful involves both rarity and probability, as demonstrated by the examples in Subsection III.C.3. Future regulatory guidance should incorporate these ideas—baseline, rarity, and probability—which are absent from the Working Party opinions on personal data and anonymization techniques.[162]

Third, applying the concept of predicate singling-out to various anonymization mechanisms highlights questions for additional study. In particular, Section IV.A identifies *k*-anonymous mechanisms that enable predicate singling-out attacks. Policymakers can clarify their understanding of singling out by referring to these and similar attacks, clarifying whether they are recognized as examples of successful GDPR singling-out attacks, and issuing guidelines accordingly. In particular, if the attacks are found legitimate in the eyes of policymakers, then they illustrate a new way that *k*-anonymity can lead to privacy failures. Additionally, the analysis of how predicate singling-out composes in Section III.E demonstrates a strong tension between allowing individual exact counts and composition. Without composition, ex-ante privacy controls are challenging.[163] Policymakers should consider whether and how ANONYMIZATION composes, and, if so, what that means for GDPR singling-out.

Fourth, predicate singling-out provides a concrete model for critique and evolution. Skeptics of our definition can revisit the myriad modeling choices involved and propose definitions that they believe better match GDPR singling-out. Identifying scenarios where these variants differ would help raise specific policy questions for comparing the definitions.[164]

Capturing these insights and the discussions above, we recommend the following model language. The language in Example 13 exemplifies how protection from singling out can be incorporated in future regulations, regulatory

[161] Linkability and inference are included together with singling out as three key criteria of effective anonymization in *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 11-12.

[162] *See Working Party Opinion on Personal Data*, *supra* note 6, at 24; *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 27, 30.

[163] *See* Fluitt et al., *supra* note 5, at 291-92.

[164] As a straw-person example, one way of modifying our model to preclude our general attacks on *k*-anonymous mechanisms is to require the attacker to succeed with probability significantly greater than 37% no matter how rare the predicate. (This is just an example; we strongly believe this variant is misguided.) That choice immediately suggests a test mechanism that is precluded by predicate singling-out but allowed by this variant. The mechanism rolls a die: if it rolls 1–4 it outputs a completely synthetic record generated totally at random; if it rolls 5 or 6, it outputs a randomly-selected record from the input dataset. We believe this mechanism enables GDPR singling-out and, therefore, is evidence against this variant of the definition.

guidance, or contractual agreements regarding rendering data anonymous or sharing anonymous data. As argued in Section III.A, predicate singling-out is a necessary condition for singling out. Thus, the recommended language describes what is necessary for protection, and may need to be extended to define sufficient protection.

Example 13: Model language for protecting against singling out

Singling out occurs when an individual person recorded in the data can be isolated (uniquely described) by a combination of specific characteristics measured in the data, and that set of characteristics is so specific as to not reasonably occur by chance in the population.

Data privacy can be reliably achieved only if the mechanism used to generate or transform data is designed protectively. It is formally impossible to assess the risk of singling out (or other privacy risks) from inspection of the data in isolation.[165] Thus, the model language requires a protective transformation to be applied before data can be considered ANONYMIZED and, thus, freed from further controls. There is potential for data privacy loss not only when data is explicitly shared, but also when data is used—even if only for aggregate statistics or decision making.[166] Thus, the language requires protections to be applied to the use of data and derivatives.

For any portion of the personal data or any information derived from the data to be considered as rendered anonymous pursuant to Recital 26 of the GDPR, a data controller must employ an appropriately protective disclosure limitation technique to limit the risk of singling out.

In practice, a substantive privacy harm can occur when the attacker increases their relative certainty unless the level of certainty is below some practical threshold limiting the action that can be taken against a data subject. Thus, the model language defines a significant increase in both relative and absolute terms.

To be considered protective, a disclosure limitation technique must guarantee that when outputs are produced by applying such technique to the data, the protected output produced by the technique cannot be used to significantly increase the likelihood of isolating individuals. An output does not significantly increase the likelihood of isolating individuals if for all combinations of characteristics measured in the data, at least one of the following is true:

[165] *See* discussion *supra* Section II.B.

[166] *See supra* note 31 and accompanying text.

> *(A) The likelihood of isolating an individual with the combination using the protected output does not have a large relative increase compared to the occurrence of the combination in the underlying population; or*
>
> *(B) The combination occurs with minimal frequency in the population and the protected output allows isolation with the combination with only minimal frequency.*
>
> Definitions for the terms "minimal frequency" and "large relative increase" should be developed based on specific policy choices regarding the appropriate tradeoff between data privacy and utility. Small thresholds for both will provide more privacy protection, at the cost of utility.[167] Appropriate thresholds will need to be selected and communicated to data controllers. We recommend that policymakers define a minimal frequency threshold that is between 0.1% and 1.0% and a large relative increase threshold that is between 5% and 20%.

B. Resolving disagreement with Article 29 Working Party Guidance

The Working Party opinion on anonymization techniques states that for *k*-anonymous mechanisms, "it should be no longer possible to single out an individual within a group of *k* users."[168] In contrast, predicate singling-out attacks are possible against a large class of *k*-anonymous mechanisms.[169] One might view this disagreement as a challenge against the argument that predicate singling-out attacks are a type of singling-out attacks.

This points to a basic question at the heart of any attempt to bridge technical and legal understanding: How should disagreements between technical findings and regulatory guidance be viewed? Three overlapping approaches are outlined below, from most deferential to the regulatory guidance to most deferential to the technical findings. These are then applied to the Working Party opinion. Ultimately, the research findings of this article are inconsistent with the Working Party opinion regarding the protection provided by *k*-anonymity. Should the EDPB find these arguments and results compelling, then it must consider adopting a position on the definition of singling out that differs from the Working Party opinion.

One alternative is that the regulatory guidance is correct by fiat—or may be made so in the future, for example as part of a legally binding decision—and without any qualification. With this approach, if there is no other way to resolve the disagreement except to alter or reject the mathematical modeling, then the mathematical modeling must give. In some cases, there may be a natural way to

[167] *See* discussion *supra* Section III.C.4.

[168] *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 16.

[169] *See supra* Section IV.A.

incorporate the regulatory guidance into an internally coherent mathematical model, perhaps by narrowing the model's applicability to scope out the point of contention. However, it may be that no coherent alternative can be found; the law is not necessarily self-consistent.

A second alternative is that the regulatory guidance is understood to apply only in typical circumstances but allows for qualification. There is an exception to every rule, and a mathematical model that admits a contrived counterexample to the legal rule need not invalidate the rule. Instead, such examples can help identify mistakes for practitioners to avoid previously unconsidered edge cases for further regulatory guidance. Updated guidance might affirm the technical findings or challenge them and leave the disagreement unresolved.

A third alternative is that the regulatory guidance is mistaken. Though it may have captured the best understanding of the issue at the time it was issued, the guidance should change in the face of evolving technology and technological understanding. It should yield to newly discovered attacks. Future regulatory guidance should reflect an updated understanding.

Whatever the approach, disagreement between regulatory guidance and a technical-mathematical analysis serves as a "stress test" for both, by uncovering and highlighting tensions. Genuine engagement with the tension will strengthen thinking on both ends of the debate. There may be a potential misunderstanding or conceptual error in the mathematical modeling, an intuitive but unrealistic expectation encoded in the regulatory guidance, etc. The disagreement can then serve as a catalyst for exchanging new insights and critiques, help guide research and modeling, and strengthen the understanding of how technological measures can meet regulatory requirements. The appropriate choice—one of the three approaches or their combination—depends on how authoritative the regulatory guidance is and how convincing the technical findings are.

Returning to the disagreement between our modeling and the Working Party on whether *k*-anonymity prevents singling out, we believe the guidance is mistaken. Namely, the use of *k*-anonymizing mechanisms should not be understood as sufficient evidence for ANONYMIZATION. The predicate singling-out attacks presented in Section IV.A apply to very large classes of *k*-anonymous mechanisms, not only very contrived examples. In particular, the attacks do not rely on mistakes in the use of *k*-anonymity, such as not considering all quasi-identifiers.[170] Moreover, the Working Party opinion is persuasive guidance rather than

[170] The Working Party opinion recognizes that poorly-used *k*-anonymity might allow GDPR singling-out. *See Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 17 ("Not considering all the quasi-identifiers when selecting the attribute to generalize is a critical mistake; if some attributes can be used to single out an individual in a cluster of *k*, then the generalization fails to protect some individuals.").

binding.[171] The opinion also recognizes that the technological understanding is evolving.[172]

It is possible that some uses of *k*-anonymity—with the right data and a mechanism applied by experts—do prevent GDPR singling-out attacks. But this requires fact-specific argument and is not the general rule.

Nevertheless, we recognize the possibility that the disagreement reflects a misunderstanding in the modeling of predicate singling-out. If so, modeling choices can be reexamined and variations that dispel the disagreement should be explored. We have considered a few variants of the definition presented in Section III.C that would exclude the attacks on *k*-anonymous mechanisms, but all had real problems. Others might find more success. One way to evaluate the possibility of misunderstanding is to consider whether predicate singling-out and *k*-anonymity mitigate certain types of informational harms that motivated the GDPR itself. If there are cases in which only predicate singling-out mitigates the harm, it could provide evidence that the model successfully captures GDPR singling-out.

VI. DISCUSSION: INTEGRATING LEGAL AND TECHNICAL REASONING FOR BETTER POLICY

A. Recommendations for hybrid legal-technical analysis

Developing hybrid legal-technical concepts is substantially more challenging than developing a concept that functions solely in one area. Based on the case of singling out above, and on multiple collaborations on legal-technical approaches to data privacy law,[173] we offer some recommendations on approaches useful in developing and evaluating hybrid legal-technical concepts.

---

[171] *See Data Protection: Opinions and Recommendations*, EUROPEAN COMM. (Nov. 24, 2016), https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/index_en.htm [https://perma.cc/4C25-W68Q] ("The material (opinions, working documents, letters etc.) issued by the Article 29 Working Party (Art. 29 WP), available on this website reflect the views only of the Art. 29 WP which has an advisory status and acts independently. They do not reflect the position of the European Commission.").

[172] *Working Party Opinion on Anonymisation Techniques*, *supra* note 6, at 4 ("[A]nonymisation and re-identification are active fields of research and new discoveries are regularly published, and on the other hand even anonymised data, like statistics, may be used to enrich existing profiles of individuals, thus creating new data protection issues.").

[173] *See generally* Kobbi Nissim & Alexandra Wood, *Is Privacy* Privacy*?*, 376 PHIL. TRANSACTIONS OF THE ROYAL SOC'Y A 20170358 1, 13-5 (2018); Micah Altman, Alexandra Wood, David R. O'Brien, Salil Vadhan & Urs Gasser, *Towards a Modern Approach to Privacy-Aware Government Data Releases*, 30 BERKELEY TECH. L.J. 1967, 2010 (2015); Kobbi Nissim, Aaron Bembenek, Alexandra Wood, Mark Bun, Marco Gaboardi, Urs Gasser, David R. O'Brien, Thomas Steinke & Salil Vadhan, *Bridging the Gap between Computer Science and Legal Approaches to Privacy*, *supra* note 90, at 763-65; Micah Altman, Alexandra Wood, David R. O'Brien & Urs Gasser, *Practical Approaches to Big Data Privacy over Time*, 8 INT'L DATA PRIV. L. 29, 39-42 (2018); Micah Altman, Stephen Chong & Alexandra Wood,

To ensure that concepts function in both realms requires functionally that both legal scholars and technical experts be able to work together to iteratively design, modify, and evaluate proposed concepts. In other words, the concept must be simultaneously definable and evaluable by legal and technical experts. To achieve this, we have found the following approaches helpful, and we conjecture that they are of broad usefulness.

The process of developing hybrid legal-technical concepts requires integration across and iteration between domains of expertise. This creates challenges for multi-disciplinary communication and evaluation. When technical or legal details are fundamentally ambiguous, even an idea considered to be unambiguous to experts from one domain may in fact be impenetrable to experts from the other. Differing assumptions by legal and technical experts over the ambiguous details can lead to divergence in analysis even when the definitions of concepts are really in agreement, or vice versa.

Simple, clearly-stated examples, including both test cases and illustrations, facilitate communication. As a demonstration, this pattern is at work throughout Parts II–IV, where we highlight instructive examples in call-out boxes. Test cases are fact patterns against which proposed concept definitions can be measured. A useful corpus of test cases will include many for which we have strong insights into the legal outcome that must result from any coherent definition of the concept, but it may also contain test cases where the outcome is less clear.

Such a corpus of test cases identifies effective technical definitions: those yielding the "right" legal answer where strong legal intuitions exist, as well as those yielding reasonable answers in cases where legal intuitions are weak. Examples may instead be illustrations of a candidate technical formalization of a concept, helping illuminate and communicate the formalization's contours. Illustrations can highlight the implications of and the assumptions underlying a formalization, allowing the candidate to be more easily evaluated by experts in either discipline. As candidates are considered and rejected, their illustrations inspire new test cases to guide future candidates.

Different domains of expertise often assign divergent definitions to the same term. Eliding these differences is a recipe for confusion. As one example, the terms anonymization and de-identification are so overloaded that productive discussion all but requires a section devoted to disambiguation.[174] Attempting to unify definitions across fields or to import concepts from one field into another can compound the confusion. For example, much of the recent wave of research on fair machine learning uses legal terms of art like "disparate treatment" and "disparate impact" to refer to specific mathematical properties inspired by their

Formalizing Privacy Laws for License Generation and Data Repository Decision Automation, arXiv:1910.10096v1 [cs.CR] 1, 3-10 (2019); Fluitt et al., *supra* note 5, at 291-92.

[174] *See supra* Section II.A.

legal antecedents.[175] Those blind to the distinction between the mathematical and legal concepts trip over this linguistic stumbling block.

A painstaking but ultimately clarifying approach is to maintain explicit parallel vocabularies of existing technical and legal terms and to clearly label each term so that its originating vocabulary is clear. For example, this article distinguishes GDPR singling-out from predicate singling-out. Rather than aim to develop a single joint definition, we aim to honor the existing definitions, indicate which we are using at different points in the analysis, and develop separate terms for hybrid concepts. Likewise, some in the algorithmic fairness community use the terms "treatment parity" and "impact parity" to distinguish from disparate treatment and impact, recognizing that "technical approaches that achieve these criteria may fail to satisfy the underlying legal and ethical desiderata."[176]

While it is unavoidable that computational experts will rely on legal scholars to validate the correctness of legal definitions (and vice versa) it is critical that experts on both sides be able to trace the consequences of key definitions in the other fields. A second, complementary approach is to develop a bilingual dictionary of sorts that provides experts in one field (such as law) with key definitions and example cases from another field (e.g., statistics) using the language of the first field.[177]

When designing hybrid concepts, be cautious of technical concepts that are defined to take on only true and false values (i.e., binary measures). While such true/false distinctions are an important part of formal reasoning, in practice, evaluation of a simple binary concept is often complicated, and yields a measure of degree. For example, creating bright-line rules, or engineering processes, based on a binary concept often requires measuring critical parameters, thresholds and conditions in a way that is dependent on the application context. More than that, identifying parameters may enable technical coherence. The case of predicate singling-out exemplifies both roles. An anonymization mechanism is secure against predicate singling-out attacks if it guarantees only an *insignificant increase* (over a statistical baseline) in the risk that an attacker can isolate a person when using *very rare* predicates. Whether an increase in risk is significant and whether a predicate is sufficiently rare are ultimately policy questions that should be informed by legal, technical, and other considerations. But our success in deriving a meaningful formalization of security against singling-out attacks hinged on the not otherwise apparent observation that it is crucial to consider the rarity of the predicate. A third parameter is made absolutely binary (and thus hidden) in the above definition: to isolate is to describe exactly one person in the dataset. But while it is clearly meaningful technically and policy-wise to also

[175] Zachary C. Lipton, Alexandra Chouldechova & Julian McAuley, *Does mitigating ML's impact disparity require treatment disparity?*, 31 ADVANCES IN NEURAL INFO. PROCESSING SYS. 8125, 8125- 8126 (2018).

[176] *Id.* at 8126.

[177] *See, e.g.*, Wood et al., *supra* note 28.

consider the case of describing a very small set of individuals, the focus on a single person is legally motivated by the GDPR.

Moving from a non-binary measure like "harm" to a bright-line threshold such as "safe" almost always involves making a policy judgement about the societally-appropriate baseline and balance among goals, among other considerations. While these judgments may be an appropriate target of legal analysis, they should be recognized as such, identified, and parameterized.[178]

For computer scientists researching legal questions, we offer the following advice based on what has worked for us so far. First, understand the legal context as well as possible. Identify authoritative texts and interpretations, as well as understand their application to specific fact patterns. Develop a corpus of examples to use as test cases and to illustrate salient features of the technical approach. Second, decide what relationship the technical approach should have with the legal question. One goal is to argue that some technical approach complies with a legal requirement, or that it does not. Another is to provide a taxonomy of relevant technical considerations, or to identify scenarios for which regulatory guidance is unclear or inconsistent. Where the gap between technical and legal thinking is very large, it is reasonable to "merely" shrink the gap, rather than attempting to build a sturdy bridge.

This article argued that a regulatory requirement entails a technical one – a test that GDPR ANONYMIZATION mechanisms had to pass. As described in Section III.B, that goal guided many of our decisions and allowed us to simplify the problem by focusing on a special case. Throughout, apply the design principle of "separation of concerns"[179] where possible. It is particularly important to distinguish between concepts that are meant to represent the high-level goals (such as fairness, social welfare, and prevention of deaths) that a body of policy is intended to advance, and concepts that represent specific characteristics or measures that are intended to align with, be instrumental for, advance, protect, or proxy those goals. Technical concepts such as predicate singling-out typically should be designed to rigorously define concepts in the latter category, and to generally (but almost never necessarily) advance concepts in the former.

Third, communicate the findings, including characterizing and describing the conditions under which a technical concept can be applied in the real world (i.e.,

[178] *See, e.g.*, Peter Abelson, *The Value of Life and Health for Public Policy*, 79 ECON. REC. S2 (2003) (analyzing technical and social choices in valuation of life years); Daniel Kifer & Ashwin Machanavajjhala, *No free lunch in data privacy*, PROC. OF THE 2011 ACM SIGMOD INT'L CONF. ON MGMT. OF DATA 193-204 (2011) (demonstrating the inevitability of value choices in privacy protection); Micah Altman, Alexandra Wood & Effy Vayena, *A Harm-Reduction Framework for Algorithmic Fairness*, 16 IEEE SEC. & PRIV. 34 (2018) (providing a framework for separating technical judgements of harm from social judgements of fairness); Effy Vayena, Urs Gasser, Alexandra Wood, David R. O'Brien & Micah Altman, *Elements of a New Ethical Framework for Big Data Research*, 72 WASH. & LEE L. REV. ONLINE 420 (2016) (providing an ethical framework for the oversight of big data research).

[179] Walter Hürsch & Cristina Videira Lopes, *Separation of Concerns*, TECHNICAL REPORT BY THE COLL. OF COMPUT. SCI., NORTHEASTERN UNIV. (1995).

how does the practitioner or scholar know when their situation fits?). This article is part of that effort.

B. Addressing the challenges of legal regulation of technically-complex domains

A careful analysis of singling out demonstrates the surprising difficulty of constructing new protections that are simultaneously practically justiciable, and technically sound. As discussed above, the development of the GDPR incorporated review by technical privacy experts which was extensive, compared to most law-making processes. Further, that review included explicit evaluation of the effectiveness of specific applied protections (such as *k*-anonymization) with respect to the GDPR's privacy goals.

Nevertheless, the extensive technical review did not expose the ambiguities of a privacy concept central to the GDPR. Moreover, the review process did not detect the fact that none of the protections under consideration could reliably achieve the desired protections. Why did these gaps escape notice? We conjecture that development of new privacy concepts that are both practically justiciable and technically coherent requires not just review by both legal and technical experts, but a co-design effort. We argue that since the process and goals of developing technical concepts and legal rules are so different the concepts developed are likely to diverge. Bringing in review at a later stage of development will be less likely to be succeed in harmonizing basic concepts than integrating development of the concepts at an earlier stage.

The traditional role of law is "as a constraint on behavior acting through the imposition of sanctions."[180] Legal concepts are developed to facilitate this role, by balancing a broad set of important but ill-defined goals, and to operate within the constraints of a legal system administered by humans with limited resources. Because of these roots, legal concepts typically reflect the complexity of human behavior, the requirements of implementation within a legal and social system, and the need for predictability in these systems.

In contrast, emerging technical approaches to privacy are heavily rooted in theoretical computer science and mathematics. A defining feature of these approaches is the use of axiomatic theory systems that precisely define axioms, concepts, symbols, operators – and from which inferences can be derived using formal logic. Within a formal framework, privacy is precisely defined, the properties of methods to protect it are rigorously derived, and it is often possible to prove that some otherwise desirable results are unachievable.

Formal approaches to privacy offer the advantage of coherence, consistency, rigor and certainty, but have the inevitable drawback of being limited in applicability. Like formal approaches to other areas of human behavior, such as those found in the fields of economics and political science, they rely on simplified conceptualizations of human goals, behavior, institutions, and societal contexts.

[180] Urs Gasser, *Recoding Privacy Law: Reflections on the Future Relationship Among Law, Technology, and Privacy*, 130 HARV. L. REV. F. 61 (2016).

Simplifications are necessary when further generalization prevents the derivation of inferences. And simplifications remain useful so long as they capture a portion of real human activity, bounded by conditions that can be reliably recognized from *outside* of the formal system.

A concept such as singling out must function effectively in both realms – as a hybrid legal-technical concept. As a legal concept, it has to be understandable to relevant actors in the legal system; it has to be achievable with feasible resources; and it has to (at least most of the time) achieve the policy goals of the law by constraining behavior that is the most harmful, while allowing productive activity in the field to which the law applies. As a technical concept, it has to be well-defined enough so that one can determine ex-ante that a specific process under examination (e.g., an algorithm or piece of software) implements the concept; one can identify (at least some of) the conditions under which it is possible to achieve the conceived goal; and one can identify logical (necessary or sufficient) relationships between this concept and meeting other goals.

Achieving either set of goals relies critically on simplifying the concept in the right way. However, simplifications made solely for legal reasons are unlikely to yield concepts that are technically tractable, and vice versa. To achieve both sets of goals at once requires co-design, to achieve a simplified version of the concept that can satisfy both legal and technical goals.

The approach to legal-technical reasoning presented in this article aims to demonstrate how systematic, rigorous analysis can be used to develop such hybrid legal-technical concepts. The result is not to replace legal or technical definitions, but to provide a framework for coordinating the two.