Epistemic Privacy

2

ALEXANDRE EVFIMIEVSKI, RONALD FAGIN, AND DAVID WOODRUFF

IBM Almaden Research Center

Abstract. We present a novel definition of privacy in the framework of offline (retroactive) database query auditing. Given information about the database, a description of sensitive data, and assumptions about users' prior knowledge, our goal is to determine if answering a past user's query could have led to a privacy breach. According to our definition, an audited property A is private, given the disclosure of property B, if no user can gain confidence in A by learning B, subject to prior knowledge constraints. Privacy is not violated if the disclosure of B causes a loss of confidence in A. The new notion of privacy is formalized using the well-known semantics for reasoning about knowledge, where logical properties correspond to sets of possible worlds (databases) that satisfy these properties. Database users are modeled as either possibilistic agents whose knowledge is a set of possible worlds, or as probabilistic agents whose knowledge is a probability distribution on possible worlds.

We analyze the new privacy notion, show its relationship with the conventional approach, and derive criteria that allow the auditor to test privacy efficiently in some important cases. In particular, we prove characterization theorems for the possibilistic case, and study in depth the probabilistic case under the assumption that all database records are considered a-priori independent by the user, as well as under more relaxed (or absent) prior-knowledge assumptions. In the probabilistic case we show that for certain families of distributions there is no efficient algorithm to test whether an audited property A is private given the disclosure of a property B, assuming $P \neq NP$. Nevertheless, for many interesting families, such as the family of product distributions, we obtain algorithms that are efficient both in theory and in practice.

Categories and Subject Descriptors: F.2.1 [**Analysis of Algorithms and Problem Complexity**]: Numerical Algorithms and Problems; H.2.0 [**Database Management**]: General—*security, integrity, and protection*; H.2.7 [**Database Management**]: Database Administration

General Terms: Algorithms, Security, Theory

Additional Key Words and Phrases: Auditing, disclosure, privacy, query logs, reasoning about knowledge, supermodularity, Positivstellensatz

ACM Reference Format:

Evfimievski, A., Fagin, R., and Woodruff, D. 2010. Epistemic privacy. J. ACM 58, 1, Article 2, (December 2010), 45 pages.
DOI = 10.1145/1870103.1870105 http://doi.acm.org/10.1145/1870103.1870105

A short version of this article is available in the *Proceedings of the 27th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS'08)*, ACM, pp. 171–180.

Authors' address: A. Evfimievski, R. Fagin, and D. Woodruff, IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120-6099, contact e-mail: evfimi@us.ibm.com.

DOI 10.1145/1870103.1870105 http://doi.acm.org/10.1145/1870103.1870105

1. *Introduction*

Today, privacy protection has become a popular and even fashionable area of database research. This situation is, of course, quite natural, given the importance of privacy in our social life and the risks we face in the digital world. These risks were highlighted by numerous recent reports of personal data theft and misappropriation, prompting many countries to enact data protection laws [Australia 1998; Canada 2000; U.S. Congress 1996; E.U. Parliament 1995]. However, the current state of scientific knowledge still does not allow the implementation of a comprehensive privacy solution that guarantees provable protection. In fact, the notion of privacy itself has many definitions and interpretations, some focused on theoretical soundness, others on practical usefulness. This article attempts to reduce the gap between these two aspects by exploring more flexible yet sound definitions.

One typical privacy enforcement problem, called *query auditing*, is to determine if answering a user's database query could lead to a privacy breach. To state the problem more accurately, we assume that the auditor is given:

—The database at the time of the user's query, or some partial knowledge about that database;
—A description of information considered sensitive, often called the *privacy policy* or the *audit query*;
—Assumptions about the user's prior knowledge of the database, of the audit query/privacy policy, and of the auditor's privacy enforcement strategy if it exists;
—The user's query, or a range of queries.

The auditor wants to check whether answering a given query could augment the user's knowledge about some sensitive data, thereby violating the privacy of that data. This problem has two extensions: *proactive* privacy enforcement by means of *online auditing* [Kenthapadi et al. 2005], and *retroactive* (*offline*) auditing.[1]

In the proactive (online) privacy enforcement scenario, users issue a stream of queries, and the database system decides whether to answer or to deny each query. The denial, when it occurs, is also an "answer" to some (implicit) query that depends on the auditor's privacy enforcement strategy, and therefore it may disclose sensitive data. The strategy has to be chosen in advance, before the user's queries become available. A strategy that protects privacy for a specified range of

[1] The term "offline" in data privacy auditing is ambiguous. We suggest to separate the following three concepts:

Online proactive auditing. The auditor receives the user's queries one by one, and must decide on the spot whether to answer the query or to refuse it, before the user submits the next query. Both the user and the auditor know the past queries and answers before they choose their next query or response.

Offline proactive auditing. The auditor receives all queries at once, and then decides which queries to answer and which to refuse. The user still learns which queries were refused, and may take advantage of this. But the user cannot choose future queries based on the past responses.

Offline retroactive auditing. The database receives queries one by one and answers all of them; no refusals. Much later, the auditor goes back and identifies all queries that could have leaked sensitive information. The user never learns which queries were identified as leaky (unless at criminal proceedings).

The last concept, offline retroactive auditing, is the one studied in this article.

queries represents a solution to this auditing problem. An in-depth discussion of online auditing can be found in Kenthapadi et al. [2005] and Nabar et al. [2006] and papers referenced therein.

In the retroactive (offline) scenario, the users issue their queries and receive the answers; later, an auditor checks if a privacy violation might have occurred. The audit results are not made available to the users, so the auditor's behavior no longer factors into the disclosure of data, and this considerably simplifies the problem. This also allows for more flexibility in defining sensitive information: while in the proactive case the privacy policy is typically fixed and open to the users, in the retroactive case the audit query itself may be sensitive, for example, based on an actual or suspected privacy breach [Agrawal et al. 2004; Motwani et al. 2008]. Retroactive auditing is the application that motivates this article, although our framework turns out to be fairly general.

To further illustrate this, suppose Alice asks Bob for his HIV status. Assume that Bob never lies and considers "HIV-positive" to be sensitive information, while "HIV-negative" is for him OK to disclose. Bob is HIV-negative at the moment; can he adopt the proactive strategy of answering "I am HIV-negative" as long as it is true? Unfortunately, this is not a safe strategy, because if he does become HIV-positive in the future, he will have to deny further inquiries, and Alice will infer that he contracted HIV. The safest bet for Bob is to always refuse an answer.[2]

For the retroactive scenario, suppose that Bob contracted HIV in 2006. Alice, Cindy and Mallory legitimately gained access to Bob's health records and learned his HIV status, but Alice and Cindy did it in 2005 and Mallory did in 2007. Bob discovers that his disease is known to the drug advertisers, and he initiates an audit, specifying "HIV-positive" as the audit query. The audit will place the suspicion on Mallory, but not on Alice and Cindy.

In legal practice, retroactive law enforcement has shown to be better suited to the complex needs of our society, although proactive measures are used too, especially in simple or critical situations. For example, a valuable item can be protected from theft by lock and key (a proactive measure) or by the fear of being caught and jailed (a retroactive measure). If it is simple to fence off the item and distribute the keys to all authorized users, or if the item has extraordinary value, then proactive defense is the best option, but in less clear-cut cases this would be too cumbersome or intrusive. After all, even an authorized user might steal or lose the item, and even a stranger sometimes should be able to gain access to it, for example, in an emergency. Healthcare [Agrawal et al. 2002] is one area where the complexity of data management is just too high to hope for a fully proactive solution to privacy. The importance of retroactive disclosure auditing in healthcare has been recognized by the U.S. President's Information Technology Advisory Committee [PITAC 2004], which recommended that healthcare information systems have the capability to audit who has accessed patient records. We believe in coexistence and importance of both auditing approaches.

1.1. Privacy Definitions in Query Auditing.

The art of encryption and cryptanalysis goes back to antiquity, but the scientific maturity of privacy theory

[2] If Alice pays Bob for answers, he can balance privacy and profit by tossing a coin and answering "I am HIV-negative" only if the coin falls heads.

was made possible only in modern times by mathematical modeling of the eavesdropper's knowledge. One of the first such models was proposed by Shannon [1949], who introduced the notion of *perfect secrecy*. Shannon suggested to represent the adversarial knowledge by a probability distribution over possible private data values: prior distribution before the cryptogram is revealed, and posterior distribution after the adversary sees the cryptogram (but not the key). Perfect secrecy corresponds to the situation where the posterior distribution is identical to the prior, for every possible cryptogram. This general idea has been later adapted and extended to many privacy frameworks and problems, including query auditing.

Denote by Ω the set of all possible databases, and by A and B two properties of these databases; each database $\omega \in \Omega$ either has or does not have each property. Assume that the actual database satisfies both A and B. Suppose that property A is sensitive, and property B is what user Alice has learned by receiving the answer to her query. Was the privacy of A violated by the disclosure of B? This depends on what Alice knew before learning B; for example, if she knew "$B \Rightarrow A$" (but did not know A), then B of course revealed to her that A is true. On the other hand, if Alice already knew that A is true, then B could no longer reveal A and may be waved through by the auditor.

Miklau and Suciu [2004] applied Shannon's model to this problem and declared A to be private given B if and only if, for all probability distributions P over Ω that might describe Alice's prior knowledge about the database, we have

$$P[A \mid B] = P[A]. \tag{1}$$

Unfortunately, if no constraints are placed on P, no pair (A, B) of nontrivial properties ($A, B \neq \varnothing$ or Ω) will satisfy this privacy definition. To see this, take a database $\omega_1 \in \Omega - B$, then take another database $\omega_2 \in \Omega$ so that $\omega_1 \in A \Leftrightarrow \omega_2 \notin A$. This is possible since neither A nor B equals $\varnothing$ or Ω. Assign the probability $P(\omega_1) = P(\omega_2) = 1/2$ and $P(\omega) = 0$ everywhere else; if the actual database, which must have a nonzero probability (see why in Remark 2.3), is $\omega^* \notin \{\omega_1, \omega_2\}$, assign $P(\omega_1) = P(\omega_2) = P(\omega^*) = 1/3$. We have $P[A \mid B] \neq P[A]$, because the (prime) denominator in $P[A]$ cannot appear in $P[A \mid B]$.

Miklau and Suciu [2004] considered a quite limiting, yet popular, constraint: that Alice treats all database records $r \in \omega$ independently, that is, P is a product distribution:

$$P(\omega) \;=\; \textstyle\prod_{r \in \omega} P[r] \;\times\; \prod_{r \notin \omega}(1 - P[r]).$$

Under this constraint, they prove that property A is private given the disclosure of B if and only if they share no *critical records* (Theorem 3.5 in Miklau and Suciu [2004]). A database record is called "critical" for A (for B) if its presence or absence in some database may decide the truth value of A (of B). This can be a real record r that belongs to the actual database ($r \in \omega^*$), or an imaginary record $r \notin \omega^*$ made up from an arbitrary combination of attribute values. For many properties A and B that, in practice, have nothing to do with each other, we can make up an imaginary record r and a pair ω_A and ω_B of imaginary databases such that inserting r into ω_A (into ω_B) flips the truth value of A (of B).

For example, if

$$\begin{aligned} A &\leftrightharpoons \quad \exists X\ \textsc{PatientID}(\text{Bob}, X)\ \&\ \textsc{Disease}(X, \text{HIV+}) \\ B &\leftrightharpoons \neg\exists Y\ \textsc{PatientID}(\text{Chris}, Y)\ \&\ \textsc{Disease}(Y, \text{HIV+}) \\ \omega^* &= \{\textsc{PatientID}(\text{Diana}, 123),\ \textsc{Disease}(123, \text{Flu})\} \end{aligned}$$

then A and B share a critical record $r = \textsc{Disease}(123, \text{HIV+})$ even though patient #123 is Diana, all patients are HIV-negative, and Bob is not even registered at the hospital. The imaginary databases are

$$\omega_A = \{\textsc{PatientID}(\text{Bob}, 123)\}; \quad \omega_B = \{\textsc{PatientID}(\text{Chris}, 123)\}.$$

One can see that, even with prior knowledge restricted to product distributions, very few practical queries would get privacy clearance: perfect secrecy appears too demanding to be practical.

A number of recent papers studied ways to relax condition (1) and make it approximate. They follow the same principle: for certain pairs (ρ_1, ρ_2) of numerical bounds, $\rho_1 < \rho_2$, require that

$$P[A] \leqslant \rho_1 \ \Rightarrow\ P[A \mid B] \leqslant \rho_2,$$

where P is a prior knowledge distribution. This idea is behind the definition of ρ_1-to-ρ_2 privacy breach in Evfimievski et al. [2003]; Kenthapadi et al. [2005] use a slightly different version as part of their definition:

$$1 - \lambda \ \leqslant\ P[A \mid B] \,/\, P[A] \ \leqslant\ 1/(1 - \lambda)$$

The Sub-Linear Queries (SuLQ) framework developed in Blum et al. [2005], Dinur and Nissim [2003], and Dwork and Nissim [2004] has a more sophisticated version with nice theoretical characteristics:

$$\mathbf{Pr}\left[\log \frac{P[A \mid B]}{1 - P[A \mid B]} - \log \frac{P[A]}{1 - P[A]} > \varepsilon\right] \leqslant \delta. \tag{2}$$

Conceptually, they all require that no user can gain much confidence in the audited property A by learning the disclosed property B, subject to prior knowledge constraints.

Perhaps surprisingly, however, all papers known to us, in their proofs if not in their definitions, do not make any distinction between *gaining* and *losing* the confidence in A upon learning B. For example, the SuLQ results remain in force if the privacy definition of Blum et al. [2005] is changed by placing the absolute value sign "|...|" over the difference in (2). In Dwork and Nissim [2004], the "|...|" appears in the definition explicitly.

It turns out that taking advantage of the gain-vs.-loss distinction yields a remarkable increase in the flexibility of query auditing. To bring it into focus, we shall put aside the approximate privacy relaxations and replace Eq. (1) with inequality

$$P[A \mid B] \ \leqslant\ P[A] \tag{3}$$

That is, we call property A *private* given the disclosure of property B when (3) holds for all distributions P that are admissible as a user's prior knowledge. One might call this "semiperfect secrecy," for it has the same sort of "absolute" form as perfect secrecy. This and related notions are the subject of this article.

Let us illustrate its flexibility with a simple example of Alice (a user) and Bob (a patient). The hospital's database ω has two records: r_1 = "Bob is HIV-positive" and r_2 = "Bob had blood transfusions." The sensitive property A is the presence of r_1, that is, the fact that Bob is HIV-positive. The property B that Alice queries and learns is "$r_1 \in \omega$ implies $r_2 \in \omega$," in other words, that "if Bob is HIV-positive, then he had blood transfusions." We make no constraints on Alice's prior knowledge distribution, other than a nonzero probability of the actual database. Could the disclosure of B violate the privacy of A? Look at the following table of possible worlds:

	$r_2 \in \omega$	$r_2 \notin \omega$
$r_1 \in \omega$	A is true	A is true ★
$r_1 \notin \omega$	A is false	A is false

For Alice, learning B has the effect of ruling out the cell marked with a ★, while leaving the other cells untouched. Whatever the cells' prior probabilities are, the odds of A can only go down: $P[A \mid B] \leqslant P[A]$. Thus, A is private with respect to B, even though A and B share a critical record r_1, and regardless of any possible dependence among the records.[3]

A closely related phenomenon was noticed in the 1940's by the mathematician George Pólya in the context of his studies of how mathematicians solve their problems. He wrote a popular and highly acclaimed book, recently reissued, about problem solving [Pólya 1957], followed by more in-depth monographs Pólya [1954, 1968]. Pólya observed the following rule of *plausible reasoning*:

$$\frac{\text{If } A \text{ then } B \quad B \text{ is true}}{A \text{ more credible}},$$

where "more credible" means that $P[A \mid B] \geqslant P[A]$. It is easy to show in the same manner that the rule holds regardless of one's prior knowledge.

1.2. SUMMARY OF RESULTS. This article studies a notion of database privacy that makes it illegal for users to gain confidence about sensitive facts, yet allows arbitrary confidence loss. We begin in Sections 2 and 3 by introducing two novel privacy frameworks that implement the above concept for two different knowledge representations: possibilistic and probabilistic. We outline some properties of our privacy definitions that are relevant to the problem of testing privacy, and give necessary and sufficient conditions for privacy with no restrictions on the user's prior knowledge.

Section 4 delves deeper into the possibilistic model. For certain important cases, notably when the constraints on a user's prior knowledge are intersection-closed (i.e., not violated by a collusion of users), we give necessary and sufficient criteria for testing possibilistic privacy, which also reduce the complexity of this problem.

Sections 5 and 6 focus on the more complex probabilistic model, over the set $\{0, 1\}^n$ of Boolean vectors that represent subsets of database records. Section 5 studies two probabilistic prior knowledge constraints: bit-wise independence (product distributions) and log-supermodularity. The bit-wise independence constraint was

[3] Note that if Bob proactively tells Alice "If I am HIV-positive, then I had blood transfusions," a privacy breach of A may occur, because Alice may learn more than just B. For example, Alice then learns that Bob is thinking about his HIV status.

used also in Miklau and Suciu [2004], so our work can be viewed as an extension of theirs. Log-supermodularity is chosen to provide a "middle ground" between bit-wise independence and the unconstrained prior knowledge. We give simple combinatorial necessary criteria and sufficient criteria for privacy under the log-supermodular and the product distribution constraints.

In Section 6, we study more general families Π of distributions over $\{0, 1\}^n$ that can be described by the intersection of a finite number of polynomial inequalities in a finite number of real-valued variables. We prove that even for certain very restricted Π, deciding whether a set $B \subseteq \{0, 1\}^n$ violates the privacy of a set $A \subseteq \{0, 1\}^n$ with respect to distributions in Π cannot be done in polynomial time, unless $P = NP$.

We overcome this negative result in two ways. First, using some deep results from algebraic geometry, we show that in certain interesting cases, such as when Π is the family of product distributions, there are provably efficient algorithms for deciding privacy. Second, we describe the sum-of-squares heuristic, introduced in Shor [1987], Shor and Stetsyuk [1997], and Parrilo [2000], and its application for deciding privacy for any Π. The heuristic has been implemented and works remarkably well in practice [Parrilo and Sturmfels 2001].

2. *Worlds and Agents*

Epistemology, the study of knowledge, has a long and honorable tradition in philosophy, starting with the early Greek philosophers. Philosophers were concerned with questions such as "What does it mean to say that someone knows something?" In the 1950's and 1960's [Hintikka 1962; Kripke 1963; van Wright 1951], the focus shifted more to developing an *epistemic logic*, a logic of knowledge, and trying to capture the inherent properties of knowledge. Here there is a set Ω of possible worlds, one of which is the "real world" ω^*. An agent's *knowledge* is a set $S \subseteq \Omega$ of worlds that the agent considers possible. Since we are modeling *knowledge* rather than *belief*, we require that $\omega^* \in S$. If F is a (possible) fact, and $A \subseteq \Omega$ is the set of possible worlds where F is true, then we say that the agent *knows* F if and only if $S \subseteq A$.

More recently, researchers in such diverse fields as economics, linguistics, artificial intelligence, and theoretical computer science have become interested in reasoning about knowledge [Fagin et al. 1995]. The focus of attention has shifted to pragmatic concerns about the relationship between knowledge and action. That is our focus: the effect of an action, such as the disclosure of certain information, on the knowledge of an agent.

Worlds. Let Ω be a finite set of all possible databases. We shall call a database $\omega \in \Omega$ a *world*, and the entire Ω *the set of all possible worlds*. The *actual world*, denoted by ω^*, represents the real database. Every property of the database, or assertion about its contents, can be formulated as "$\omega^* \in A$" where $A \subseteq \Omega$ is the set of all databases that satisfy the property. A subset $A \subseteq \Omega$ that contains ω^* shall be called a *knowledge set*.

Agents. We shall think of database users as *agents* who know something about the worlds in Ω and who try to figure out which $\omega \in \Omega$ is the actual world ω^*. An agent's knowledge can be modelled in different ways; we shall consider two

approaches. In a *possibilistic* agent, knowledge is represented by a set $S \subseteq \Omega$ that contains exactly all the worlds this agent considers possible. In particular, $\omega^* \in S$. Here every world is either possible or not, with no ranking or score assigned. In a *probabilistic* agent, knowledge is represented by a probability distribution $P : \Omega \to \mathbb{R}_+$ that assigns a nonnegative weight $P(\omega)$ to every world. We denote the sum $\sum_{\omega \in A} P(\omega)$ by $P[A]$, requiring that $P[\Omega] = 1$ and $P(\omega^*) > 0$; by $\mathbb{R}_+$ we denote the set of all non-negative real numbers.

We say that a possibilistic agent with knowledge S *knows* a property $A \subseteq \Omega$ when $S \subseteq A$. We say that A is *possible* for this agent when $S \cap A \neq \varnothing$, that is, when the agent does not know $\Omega - A$. For a probabilistic agent with distribution P, to *know* A means to have $P[A] = 1$, and to consider A possible means to have $P[A] > 0$.

A function Q whose domain is Ω shall be called a *query*; if its range is $\{0, 1\}$ then Q is a *Boolean* query. For a given actual world ω^*, each query Q corresponds to the knowledge set associated with the query's "actual" output: $\{\omega \in \Omega \mid Q(\omega) = Q(\omega^*)\}$.

The Auditor. There is a special "meta-agent" called *the auditor* whose task is to analyze the queries disclosed to the users and determine which of these disclosures could breach privacy. The auditor may or may not have complete information about the actual world ω^*. For example, if the query disclosure occurred several years ago, the record update logs may provide only a partial description of the database state at that moment. Even more importantly, the auditor does not know what the user's knowledge of the database was at the disclosure time. We characterize the auditor's knowledge by specifying which pairs of a database ω and the user's knowledge S (or P) the auditor considers possible. Let us formally define the auditor's knowledge about a user:

Definition 2.1. (Possibilistic case) A *possibilistic knowledge world* is a pair (ω, S), where ω is a world and S is a knowledge set, which satisfies $\omega \in S \subseteq \Omega$. The set of all possibilistic knowledge worlds shall be denoted as

$$\Omega_{\text{poss}} := \{(\omega, S) \mid \omega \in S \subseteq \Omega\}.$$

Ω_{poss} can be viewed as an extension of Ω. For a given user whose knowledge is $S^* \subseteq \Omega$, the pair $(\omega^*, S^*) \in \Omega_{\text{poss}}$ is called the *actual* knowledge world. The auditor's knowledge about the user is defined as a nonempty set $K \subseteq \Omega_{\text{poss}}$ of knowledge worlds, which must include the actual knowledge world. We refer to K as a *second-level knowledge set*.

We now give the intuition behind a second-level knowledge set K. Assume $K = \{(\omega_1, S_1), (\omega_2, S_2), \ldots\}$. Then, the auditor knows that either (i) ω_1 is the actual world and the agent's knowledge set is S_1 (the latter means that the agent knows that the actual world is contained in S_1), or (ii) ω_2 is the actual world and the agent's knowledge set is S_2, or . . . . In particular, the auditor knows that (a) the actual world is one of $\omega_1, \omega_2, \ldots$, and the auditor knows that (b) the agent's knowledge set is one of $S_1, S_2, \ldots$. The second-level knowledge set provides richer knowledge for the auditor than simply the knowledge of (a) and (b) together, since the second-level knowledge set ties together choices for the actual world with choices for the agent's knowledge set. Note also that if the auditor knows that the actual world is ω^*, then the second-level knowledge set is of the form $\{(\omega^*, S_1), (\omega^*, S_2), \ldots\}$.

Our knowledge worlds (ω, S) are similar to the 2-worlds of Fagin et al. [1991], except that the 2-worlds of Fagin et al. [1991] would deal not only with the knowledge that the user has of the world, but also with the knowledge that the auditor has of the world. Also, our second-level knowledge sets are similar to the 3-worlds of Fagin et al. [1991], except that the 3-worlds of Fagin et al. [1991] would deal not only with the knowledge that the auditor has about the user's knowledge of the world, but also with the knowledge that the user has about the auditor's knowledge of the world.

Definition 2.2. (Probabilistic case) A *probabilistic knowledge world* is a pair (ω, P) where P is a probability distribution over Ω such that $P(\omega) > 0$. The set of all probabilistic knowledge worlds shall be denoted as

$$\Omega_{\text{prob}} := \{(\omega, P) \mid P \text{ is a distribution, } P(\omega) > 0\}.$$

The actual knowledge world $(\omega^*, P^*) \in \Omega_{\text{prob}}$ and the auditor's second-level knowledge set $K \subseteq \Omega_{\text{prob}}$ are defined analogously to the possibilistic case.

Remark 2.3. The requirement of $\omega \in S$ for every pair $(\omega, S) \in \Omega_{\text{poss}}$ and of $P(\omega) > 0$ for every pair $(\omega, P) \in \Omega_{\text{prob}}$ represent our assumption that every agent considers the actual world possible. All pairs that violate this assumption are excluded as inconsistent. Note that a probabilistic pair (ω, P) is consistent if and only if the possibilistic pair $(\omega, \text{supp}(P))$ is consistent, where $\text{supp}(P)$ is defined next.

Definition 2.4. The *support set* of a probability distribution P over Ω is the set $\text{supp}(P) := \{\omega \mid P(\omega) > 0\}$. For a family Π of probability distributions over Ω, we define a family $\text{supp}(\Pi)$ of nonempty subsets of Ω as follows: $\text{supp}(\Pi) := \{\text{supp}(P) \mid P \in \Pi\}$.

Remark 2.5. In practice, it may be computationally infeasible to precisely characterize the auditor's second-level knowledge and to use this precisely characterized knowledge in the privacy definitions. Instead, the auditor makes assumptions about the database and the user's knowledge by placing constraints on the possible pairs (ω, S) or (ω, P). These assumptions and constraints are also represented by a second-level knowledge set, which must contain the auditor's precise knowledge set as a subset. From now on, when we talk about the auditor's knowledge set, we mean the assumptions, accepted by the auditor, that form a superset of the actual knowledge set, unless stated otherwise.

Definitions 2.1 and 2.2 allow us to consider an auditor whose assumptions about the user's knowledge depend on the contents of the database. For example, the auditor may assume that, if the hospital database contains record "Bob's doctor is Alice," then Alice knows Bob's HIV status, but if there is no such record, then Alice may or may not know it. However, in many situations we can separate the auditor's knowledge about the database from the auditor's assumptions about the user. We do so by specifying two sets:

(1) A nonempty set $C \subseteq \Omega$ that consists of all databases the auditor considers possible, with $\omega^* \in C$;
(2) A family Σ of subsets of Ω and/or a family Π of probability distributions over Ω. The possibilistic agent's knowledge has to belong to Σ, and the probabilistic agent's knowledge has to belong to Π.

If the auditor knows the actual database exactly, for example, by reconstructing its state from the update logs, then $C = \{\omega^*\}$; if the auditor has no information about the database or is unwilling to take advantage of it, then $C = \Omega$. Some choices for Σ and Π will be discussed in the subsequent sections.

When we say that the auditor's knowledge is represented by C and Σ described above, we mean that all knowledge worlds (ω, S) with $\omega \in C$ and $S \in \Sigma$, and none other, are considered possible by the auditor. However, in most cases the auditor's second-level knowledge set cannot be the Cartesian product $C \times \Sigma$, because it contains inconsistent (ω, S) pairs (see Remark 2.3). The same is true in the probabilistic case, for C and Π. Let us then define a product operation that excludes all inconsistent pairs:

Definition 2.6. The *product* of a set $C \subseteq \Omega$ and a family Σ of subsets of Ω (a family Π of probability distributions over Ω) is a second-level knowledge set $C \otimes \Sigma$ ($C \otimes \Pi$) defined by

$$C \otimes \Sigma := \{(\omega, S) \in C \times \Sigma \mid \omega \in S\} = (C \times \Sigma) \cap \Omega_{\text{poss}}$$
$$C \otimes \Pi := \{(\omega, P) \in C \times \Pi \mid P(\omega) > 0\} = (C \times \Pi) \cap \Omega_{\text{prob}}$$

We call the pair (C, Σ) or (C, Π) *consistent* if their product $C \otimes \Sigma$ or $C \otimes \Pi$ is nonempty, because $\varnothing$ is not a valid second-level knowledge set.

Remark 2.7. The product $C \otimes \Sigma$ (or $C \otimes \Pi$) computes the *maximum* second-level knowledge set $K \subseteq \Omega_{\text{poss}}$ (or $K \subseteq \Omega_{\text{prob}}$) that is a subset of $C \times \Sigma$ (or $C \times \Pi$).

The auditor can safely discard from Σ all sets that have empty intersection with C, and from Π all probabilities P that have $P[C] = 0$, because they do not allow $\omega^* \in C$ as a possibility. In particular, the empty set $\varnothing$, if present in Σ, is always discarded.[4] In the same way, a world $\omega \in C$ can be safely discarded if for all $S \in \Sigma$ ($P \in \Pi$) we have $\omega \notin S$ ($P(\omega) = 0$). When a pair (C, Σ) has nothing to discard in this manner, we shall call it *nonexcessive*; analogously for (C, Π).

Remark 2.8. It is easy to see that the following conditions are equivalent:

(1) Pair (C, Σ) is nonexcessive;
(2) $\pi_1(C \otimes \Sigma) = C$ and $\pi_2(C \otimes \Sigma) = \Sigma$, where π_i is the projection operation;
(3) $\exists\, K \subseteq \Omega_{\text{poss}}$ such that $C = \pi_1(K)$ and $\Sigma = \pi_2(K)$;
(4) In the bipartite graph with vertices $\omega \in C$ and $S \in \Sigma$, where (ω, S) is an edge if and only if $\omega \in S$, there are no isolated vertices.

A probabilistic-knowledge pair (C, Π) is nonexcessive if and only if the possibilistic knowledge pair $\big(C,\ \text{supp}(\Pi)\big)$ is nonexcessive.

3. *Privacy of Knowledge*

This section introduces the definition of privacy for the possibilistic and the probabilistic knowledge models. Let $A, B \subseteq \Omega$ be two arbitrary nonempty subsets of Ω;

[4] The empty set may be added to Σ in order to make it $\cap$-closed: $\forall S_1, S_2 \in \Sigma : S_1 \cap S_2 \in \Sigma$. See Section 4.1 for more on $\cap$-closed knowledge.

as a shorthand, write $\bar{A} = \Omega - A$ and $AB = A \cap B$. Sets A and B correspond to two Boolean queries on the database ω^*; for example, query A returns "true" if $\omega^* \in A$ and "false" otherwise.

We shall study the following question: When could the disclosure of B violate the privacy of A? In our model, a positive result of query A is considered private and needs protection, whereas a negative result (that asserts $\bar{A}$) is not protected. Neither the user nor the auditor are assumed to know if A is true, and A may actually be false. On the other hand, B represents the disclosed fact, and therefore B has to be true. The auditor knows that B is true; the user transitions from not knowing B to knowing B.

The user modifies his knowledge when he receives a disclosed query result. The disclosed knowledge set $B \subseteq \Omega$ tells him that every world in $\Omega - B$ is impossible. We model the user's acquisition of B as follows. A possibilistic agent with prior knowledge $S \subseteq \Omega$, upon receiving B such that $SB \neq \varnothing$ (because $\omega^* \in SB$), ends up with posterior knowledge SB. A probabilistic agent with prior distribution $P : \Omega \to \mathbb{R}_+$, upon receiving B such that $P[B] \geqslant P(\omega^*) > 0$, ends up with posterior distribution $P(\cdot \mid B)$ defined by

$$P(\omega \mid B) = \begin{cases} P(\omega)/P[B], & \omega \in B \\ 0, & \omega \in \Omega - B \end{cases}$$

Notice that the acquisition of B_1 followed by B_2 is equivalent to the acquisition of $B_1 B_2 = B_1 \cap B_2$.

Conceptually, we say that property A is private, given the disclosure of property B, if the user could not gain confidence in A by learning B. Below, we shall make this notion precise for the two knowledge models, possibilistic and probabilistic. From this section on, we shall use pronoun "he" for the user and "she" for the auditor.

3.1. POSSIBILISTIC PRIVACY. Let us suppose first that the auditor knows everything: the actual database ω^* such that $\omega^* \in B$, and the actual knowledge set S^* of the user at the time of the disclosure. In the possibilistic model, the user may have only two "grades of confidence" in property A: he either knows A ($S^* \subseteq A$), or he does not ($S^* \not\subseteq A$). The user gains confidence when he does not know A before learning B (i.e. $S^* \not\subseteq A$) and knows A after learning B (i.e. $S^* \cap B \subseteq A$). Therefore, the privacy of A is preserved if and only if $\neg\,(S^* \not\subseteq A \;\&\; S^* \cap B \subseteq A)$, or equivalently, if and only if

$$S^* \cap B \subseteq A \;\;\Rightarrow\;\; S^* \subseteq A. \tag{4}$$

Now, suppose that the auditor does not know ω^* and S^* precisely, but has a second-level knowledge set $K \subseteq \Omega_{\text{poss}}$ such that $(\omega^*, S^*) \in K$. Then, the auditor makes sure that A is private given B by checking condition (4) for all pairs in K. Before doing so, the auditor must discard from K all pairs (ω, S) such that $\omega \notin B$, because they are inconsistent with the disclosure of B. We arrive at the following possibilistic privacy definition:

Definition 3.1. Set $A \subseteq \Omega$ is called *K-private* given the disclosure of set $B \subseteq \Omega$, for $K \subseteq \Omega_{\text{poss}}$, when

$$\forall\,(\omega, S) \in K : \;\; (\omega \in B \;\;\&\;\; S \cap B \subseteq A) \;\;\Rightarrow\;\; S \subseteq A. \tag{5}$$

We denote this predicate by $\text{Safe}_K(A, B)$.

When the auditor wants to separate her knowledge about the database from her assumptions about the user's knowledge, she represents her second-level knowledge set K as a product $C \otimes \Sigma$, where $C \subseteq \Omega$ and Σ is a family of subsets of Ω. In this case, we shall use the term "(C, Σ)-private" and the notation $\text{Safe}_{C,\Sigma}(A, B)$, which is defined as $\text{Safe}_{C\otimes\Sigma}(A, B)$. We use $\mathcal{P}(\Omega)$ to denote the power set of Ω.

PROPOSITION 3.2. *For a consistent pair* (C, Σ) *such that* $C \subseteq \Omega$ *and* $\Sigma \subseteq \mathcal{P}(\Omega)$*, the privacy predicate* $\text{Safe}_{C,\Sigma}(A, B)$ *can be equivalently defined as follows (denoting* $S \cap B \cap C$ *as* SBC*):*

$$\forall S \in \Sigma\colon\ (SBC \neq \varnothing\ \ \&\ \ SB \subseteq A)\ \Rightarrow\ S \subseteq A. \tag{6}$$

PROOF. The following sentences are trivially equivalent:

$$\begin{aligned}
&\forall S \in \Sigma\colon && (SBC \neq \varnothing\ \ \&\ \ SB \subseteq A)\ \Rightarrow\ S \subseteq A\\
&\forall S \in \Sigma\colon && (\exists\, \omega \in SC\colon\ \omega \in B\ \ \&\ \ SB \subseteq A)\ \Rightarrow\ S \subseteq A\\
&\forall S \in \Sigma,\ \forall \omega \in SC\colon && (\omega \in B\ \ \&\ \ SB \subseteq A)\ \Rightarrow\ S \subseteq A\\
&\forall (\omega, S) \in C \otimes \Sigma\colon && (\omega \in B\ \ \&\ \ SB \subseteq A)\ \Rightarrow\ S \subseteq A.
\end{aligned}$$

Thus, we have (6) $\Leftrightarrow$ (5) for $K = C \otimes \Sigma$. □

3.2. PROBABILISTIC PRIVACY. Once again, suppose first that the auditor knows the actual database $\omega^* \in B$ and the actual probability distribution P^* that represents the user's knowledge prior to the disclosure. As opposed to Section 3.1, in the probabilistic model the user has a continuum of "grades of confidence" in A, measured by $P^*[A]$. The user gains confidence whenever his *prior* probability of A before learning B, which is $P^*[A]$, is strictly smaller than his *posterior* probability of A after B is disclosed, which is $P^*[A \mid B]$. Therefore, the privacy of A is preserved if and only if

$$P^*[A \mid B]\ \leqslant\ P^*[A]. \tag{7}$$

The conditional probability $P^*[A \mid B]$ is well-defined since $P^*[B] \geqslant P^*(\omega^*) > 0$.

When the auditor does not know ω^* and P^*, but has a second-level knowledge set $K \subseteq \Omega_{\text{prob}}$ such that $(\omega^*, P^*) \in K$, she has to check inequality (7) for all possible pairs (ω, P) in K. Before doing so, she must discard all pairs (ω, P) such that $\omega \notin B$. We obtain the following probabilistic privacy definition:

Definition 3.3. Set $A \subseteq \Omega$ is called *K-private* given the disclosure of set $B \subseteq \Omega$, for $K \subseteq \Omega_{\text{prob}}$, when

$$\forall (\omega, P) \in K\ :\ \ \omega \in B\ \ \Rightarrow\ \ P[A \mid B]\ \leqslant\ P[A]. \tag{8}$$

As before, we denote this predicate by $\text{Safe}_K(A, B)$.

When the auditor's knowledge can be represented as a product $C \otimes \Pi$ for some $C \subseteq \Omega$ and some family Π of probability distributions over Ω, we shall use the term "(C, Π)-private" and the notation $\text{Safe}_{C,\Pi}(A, B)$, which is defined as $\text{Safe}_{C\otimes\Pi}(A, B)$. In this case the following proposition can be used:

PROPOSITION 3.4. *For a consistent pair* (C, Π) *where* $C \subseteq \Omega$ *and* Π *is a family of distributions over* Ω*, the privacy predicate* $\text{Safe}_{C,\Pi}(A, B)$ *can be equivalently defined as follows:*

$$\forall P \in \Pi\colon\ P[BC] > 0\ \Rightarrow\ P[AB]\ \leqslant\ P[A]\,P[B]. \tag{9}$$

PROOF. The following sentences are trivially equivalent:

$$\begin{aligned}
&\forall P \in \Pi: && P[BC] > 0 \;\Rightarrow\; \text{ineq}\\
&\forall P \in \Pi: && (\exists \omega \in BC:\ P(\omega) > 0) \;\Rightarrow\; \text{ineq}\\
&\forall P \in \Pi,\ \forall \omega \in C: && (P(\omega) > 0 \ \&\ \omega \in B) \;\Rightarrow\; \text{ineq}\\
&\forall (\omega, P) \in C \otimes \Pi: && \omega \in B \;\Rightarrow\; \text{ineq},
\end{aligned}$$

where "ineq" stands for "$P[AB] \leqslant P[A]\,P[B]$," which is equivalent to "$P[A \mid B] \leqslant P[A]$" as long as the left-hand side of the implication is true. Thus, we have (9) $\Leftrightarrow$ (8) for $K = C \otimes \Pi$. □

In fact, the definition of privacy given by (9) can be further simplified, for many families Π that occur in practice:

Definition 3.5. For a family Π of distributions over Ω, denote

$$\text{Safe}_{\Pi}(A, B) \iff \forall P \in \Pi:\ P[AB] \leqslant P[A]\,P[B]. \tag{10}$$

Notice that $\text{Safe}_{\Pi}(A, B)$ is symmetric with respect to A and B, which may not be the case for $\text{Safe}_{C,\Pi}(A, B)$. Let us state the relationship between these two predicates after the following definition:

Definition 3.6. We shall call a family Π *ω-liftable* for $\omega \in \Omega$ when $\forall P \in \Pi$ such that $P(\omega) = 0$ it satisfies the condition

$$\forall \varepsilon > 0\ \ \exists P' \in \Pi:\ P'(\omega) > 0\ \ \&\ \ ||P - P'||_{\infty} < \varepsilon. \tag{11}$$

Family Π is called *S-liftable* for a set $S \subseteq \Omega$ when Π is ω-liftable for all $\omega \in S$. The norm $||P - P'||_{\infty} := \max_{\omega \in \Omega} |P(\omega) - P'(\omega)|$.

PROPOSITION 3.7. *For every consistent pair (C, Π) and for all $A, B \subseteq \Omega$ such that $BC \neq \varnothing$ (since $\omega^* \in BC$), we have:*

$$\begin{aligned}
&\text{Safe}_{\Pi}(A, B) \;\Rightarrow\; \text{Safe}_{C,\Pi}(A, B);\\
&\text{Safe}_{C,\Pi}(A, B)\ \ \&\ \ \Pi \textit{ is } C\textit{-liftable} \;\Rightarrow\; \text{Safe}_{\Pi}(A, B).
\end{aligned} \tag{12}$$

PROOF. Trivially, the definition (10) for $\text{Safe}_{\Pi}(A, B)$ implies the characterization (9) for $\text{Safe}_{C,\Pi}(A, B)$. To prove implication (12), assume that (9) holds, but $\text{Safe}_{\Pi}(A, B)$ does not hold, and arrive at a contradiction. Take some $\omega \in BC$ and $P \in \Pi$ such that $P[AB] > P[A]\,P[B]$, to violate (10). By (9), we must have $P[BC] = 0$, so in particular $P(\omega) = 0$. However, since $\omega \in C$ and Π is C-liftable, we can use condition (11) and pick $P' \in \Pi$ that is close enough to P to still have $P'[AB] > P'[A]\,P'[B]$, yet already $P'(\omega) > 0$ and $P'[BC] > 0$, violating (9). □

3.3. PROPERTIES OF PRIVACY

Conservative Assumptions. It is easy to see from Definitions 3.1 and 3.3 that $\text{Safe}_K(A, B)$ and $K' \subseteq K$ imply $\text{Safe}_{K'}(A, B)$, in both the possibilistic and the probabilistic models. As a special case, if $C' \subseteq C$, $\Sigma' \subseteq \Sigma$, and $\Pi' \subseteq \Pi$, then $\text{Safe}_{C,\Sigma}(A, B) \Rightarrow \text{Safe}_{C',\Sigma'}(A, B)$, and $\text{Safe}_{C,\Pi}(A, B) \Rightarrow \text{Safe}_{C',\Pi'}(A, B)$. Therefore, the auditor may assume less than she actually knows (i.e., consider more knowledge worlds possible) and still catch all privacy violations, at the expense of restricting more queries.

Disclosing Less Knowledge. In the possibilistic model, for all second-level knowledge sets $K \subseteq \Omega_{\text{poss}}$ that the auditor might have, for all private properties $A \subseteq \Omega$ and for all sets $B, B' \subseteq \Omega$ we have:

$$\text{Safe}_K(A, B) \;\&\; \text{Safe}_K(A, B') \;\Rightarrow\; \text{Safe}_K(A, B \cup B'). \tag{13}$$

This immediately follows from (5) once we observe that $\omega \in B \cup B'$ implies one of $\omega \in B$ or $\omega \in B'$. Moreover, if the auditor has excluded from K all pairs (ω, S) such that $\omega \notin B$, then the condition "$\text{Safe}_K(A, B')$" is not necessary in (13): $\forall\, B, B' \subseteq \Omega$

$$\pi_1(K) \subseteq B \;\&\; \text{Safe}_K(A, B) \;\Rightarrow\; \text{Safe}_K(A, B \cup B').$$

In other words, in the possibilistic model it is always safer when less information has been disclosed. However, in the probabilistic model, even the privacy preservation under union (13) does not hold. Take, for example, $\Omega = \{1, \ldots, 6\}$, $K = \{1\} \otimes \{P\}$ where $P =$ uniform distribution, $A = \{1, 2, 3, 4\}$, $B = \{1, 2, 6\}$ and $B' = \{1, 3, 6\}$; then we have:

$$P[A] = P[A \mid B] = P[A \mid B'] = 2/3 \;<\; 3/4 = P[A \mid B \cup B'].$$

Both $\text{Safe}_K(A, B)$ and $\text{Safe}_K(A, B')$ hold, but $\text{Safe}_K(A, B \cup B')$ does not hold.

Probabilities Refine Possibilities. If $P : \Omega \to \mathbb{R}_+$ is a probability distribution that represents a user's knowledge, then its support set $\text{supp}(P)$ is the set of all worlds that this user considers possible. More generally, every second-level probabilistic knowledge set $K \subseteq \Omega_{\text{prob}}$ can be converted into the possibilistic knowledge set

$$K' \;=\; \{(\omega, \text{supp}(P)) \mid (\omega, P) \in K\}.$$

It is easy to check directly by verifying Definition 3.1 that

$$\forall\, A, B \subseteq \Omega: \quad \text{Safe}_K(A, B) \;\Rightarrow\; \text{Safe}_{K'}(A, B). \tag{14}$$

Indeed, for every (ω, S) in K' such that $\omega \in B$ and $S \cap B \subseteq A$, take $(\omega, P) \in K$ such that $S = \text{supp}(P)$. We have $P[A \mid B] = 1$ because A has all the support of P that lies inside B, and we have $P[A \mid B] \leqslant P[A]$ because $\text{Safe}_K(A, B)$, $(\omega, P) \in K$, and $\omega \in B$ (see Definition 3.3). Therefore, $P[A] = 1$ too, implying $S = \text{supp}(P) \subseteq A$.

Equation (14) gives a useful necessary condition for $\text{Safe}_K(A, B)$. Also, as we shall see in Section 5, it helps to understand K-privacy better by focusing our attention on the important aspects of the auditor's probabilistic knowledge assumption.

For the simplified privacy predicate $\text{Safe}_\Pi(A, B)$ introduced in Definition 3.5, where Π is a family of probabilities, we can make (14) slightly stronger and write, for $\Sigma = \text{supp}(\Pi)$:

$$\forall\, A, B \subseteq \Omega: \quad \text{Safe}_\Pi(A, B) \;\Rightarrow\; \text{Safe}_{\Omega,\Sigma}(A, B) \;\&\; \text{Safe}_{\Omega,\Sigma}(\bar{A}, \bar{B}). \tag{15}$$

First, $\text{Safe}_\Pi(A, B) \Rightarrow \text{Safe}_{\Omega,\Pi}(A, B)$ by Proposition 3.7, which in turn implies $\text{Safe}_{\Omega,\Sigma}(A, B)$ by (14); and second, $\text{Safe}_\Pi(A, B) \Leftrightarrow \text{Safe}_\Pi(\bar{A}, \bar{B})$ due to the following proposition:

PROPOSITION 3.8. *For all $A, B \subseteq \Omega$ and for all probability distributions P over Ω, we have:*

$$\begin{aligned} P[A]\,P[B] - P[AB] &= P[A\bar{B}]\,P[\bar{A}B] - P[AB]\,P[\bar{A}\bar{B}] \\ &= P[\bar{A}]\,P[\bar{B}] - P[\bar{A}\bar{B}]. \end{aligned} \tag{16}$$

PROOF. The first equality can be obtained as follows:

$$\begin{aligned} &P[A]\,P[B] - P[AB]\cdot 1 \\ &\quad = \big(P[AB] + P[A\bar{B}]\big)\big(P[AB] + P[\bar{A}B]\big) - P[AB]\big(P[AB] + P[\bar{A}B] \\ &\qquad + P[A\bar{B}] + P[\bar{A}\bar{B}]\big) = P[A\bar{B}]\,P[\bar{A}B] - P[AB]\,P[\bar{A}\bar{B}]. \end{aligned}$$

Equality (16) follows by symmetry. □

Multiple Disclosures. Assume that the user learns knowledge set B_1 followed by B_2, which is equivalent to the acquisition of B_1B_2. When the auditor's second-level knowledge set K represents her assumption about the user's knowledge, rather than her knowledge of the user's knowledge (see Remark 2.5), she may want to require that K remains a valid assumption after each disclosure. This property is formalized below:

Definition 3.9. Let K be a second-level knowledge set, which may be possibilistic ($K \subseteq \Omega_{\text{poss}}$) or probabilistic ($K \subseteq \Omega_{\text{prob}}$). A set $B \subseteq \Omega$ is called *K-preserving* when

Possibilistic K. For all $(\omega, S) \in K$ such that $\omega \in B$, we have $(\omega, S \cap B) \in K$;
Probabilistic K. For all $(\omega, P) \in K$ such that $\omega \in B$, we have $(\omega, P(\cdot \mid B)) \in K$.

Suppose that knowledge sets B_1 and B_2 are individually safe to disclose, while protecting the privacy of A, to an agent whose knowledge satisfies the constraints defined by K. If, after B_1 is disclosed, the updated agent's knowledge still satisfies the constraints, then it is safe to disclose B_2 too. Thus, it is safe to disclose both sets at once—as long as at least one of them preserves the constraints:

PROPOSITION 3.10. *For every second-level knowledge set K, possibilistic or probabilistic, we have:*

(1) *B_1 and B_2 are K-preserving $\Rightarrow$ B_1B_2 is K-preserving;*
(2) *If* $\text{Safe}_K(A, B_1)$ *and* $\text{Safe}_K(A, B_2)$ *and if at least one of B_1, B_2 is K-preserving, then* $\text{Safe}_K(A,\ B_1B_2)$.

PROOF

(1) trivially holds; just notice that Definition 3.9 checks knowledge worlds $(\omega, S) \in K$ or $(\omega, P) \in K$ only where both $\omega \in B_1$ and $\omega \in B_2$;

(2) Without loss of generality, assume that B_1 is K-preserving. If K is possibilistic, we must take an arbitrary $(\omega, S) \in K$ such that $\omega \in B_1B_2$ and $SB_1B_2 \subseteq A$, and show that $S \subseteq A$. Indeed, we have $(\omega, SB_1) \in K$ because B_1 is K-preserving, $SB_1 \subseteq A$ by applying K-privacy definition (Definition 3.1) to B_2, and $S \subseteq A$ by applying K-privacy definition to B_1.

If K is probabilistic, take an arbitrary $(\omega, P) \in K$ such that $\omega \in B_1 B_2$, and denote $P_1 := P(\cdot \mid B_1)$. We have $(\omega, P_1) \in K$ because B_1 is K-preserving, and

$$\begin{aligned} P[A \mid B_1 B_2] &= \frac{P[A \cap B_1 B_2]}{P[B_1 B_2]} = \frac{P[A \cap B_2 \mid B_1]}{P[B_2 \mid B_1]} \\ &= P_1[A \mid B_2] \leqslant P_1[A] := P[A \mid B_1] \leqslant P[A], \end{aligned}$$

by applying K-privacy definition (Definition 3.3) first to P_1 and B_2, then to P and B_1. □

Remark 3.11. Proposition 3.10 implies that both the family of K-preserving sets and its sub-family of the K-preserving sets safe to disclose while protecting A are $\cap$-closed. Without the "K-preserving" constraint, the family of sets that are safe to disclose does not have to be $\cap$-closed (Remark 4.2). See Section 4 and especially Theorem 4.14 for a class of situations where the "K-preserving" constraint can be lifted.

3.4. UNRESTRICTED PRIOR KNOWLEDGE. What is the characterization of privacy when the auditor knows nothing? More formally, which knowledge sets A and B satisfy K-privacy for $K = \Omega_{\text{poss}} = \Omega \otimes \mathcal{P}(\Omega)$ and for $K = \Omega_{\text{prob}} = \Omega \otimes \mathcal{P}^{\text{prob}}(\Omega)$, where $\mathcal{P}^{\text{prob}}(\Omega)$ is the set of all probability distributions over Ω? Also, what is the answer to this question if the auditor has complete information about the actual world ω^*, but knows nothing about the user's knowledge, that is, for $K = \{\omega^*\} \otimes \mathcal{P}(\Omega)$ and for $K = \{\omega^*\} \otimes \mathcal{P}^{\text{prob}}(\Omega)$? Here is a theorem that answers these questions:

THEOREM 3.12. *For all sets $A, B \subseteq \Omega$ and for all $\omega^* \in B$ the following four conditions are equivalent:*

(1) $\text{Safe}_K(A, B)$ *for* $K = \Omega_{\text{poss}}$;
(2) $\text{Safe}_K(A, B)$ *for* $K = \Omega_{\text{prob}}$;
(3) $\text{Safe}_K(A, B)$ *for* $K = \{\omega^*\} \otimes \mathcal{P}^{\text{prob}}(\Omega)$;
(4) *Either* $A \cap B = \varnothing$, *or* $A \cup B = \Omega$.

Also, the following two conditions are equivalent (again $\omega^* \in B$):

(i) $\text{Safe}_K(A, B)$ *for* $K = \{\omega^*\} \otimes \mathcal{P}(\Omega)$;
(ii) $A \cap B = \varnothing$, *or* $A \cup B = \Omega$, *or* $\omega^* \notin A$.

PROOF. First, we assume condition (4), that is, either $A \cap B = \varnothing$ or $A \cup B = \Omega$, and prove $\text{Safe}_K(A, B)$ for all second-level knowledge sets of the form $K = C \otimes \mathcal{P}(\Omega)$ and $C \otimes \mathcal{P}^{\text{prob}}(\Omega)$, including those where $C = \Omega$ or $\{\omega^*\}$. In the possibilistic case, by Proposition 3.2 it is enough to check implication (6), which is:

$$\forall S \in \Sigma\colon\ (SBC \neq \varnothing \ \&\ SB \subseteq A) \ \Rightarrow\ S \subseteq A. \tag{17}$$

If $AB = \varnothing$, then $SB \subseteq A \ \Rightarrow\ SB = \varnothing$, making the left-hand side of (17) always false and the entire implication true. If $A \cup B = \Omega$, then $S - B \subseteq A$, hence $SB \subseteq A$ alone implies $S \subseteq A$ in (17). In the probabilistic case, for every $P \in \mathcal{P}^{\text{prob}}(\Omega)$, the privacy inequality $P[AB] \leqslant P[A]\, P[B]$ trivially holds when

$AB = \varnothing$, and holds when $A \cup B = \Omega$ due to Proposition 3.8:

$$P[A]\,P[B] \;-\; P[AB] \;=\; P[\bar{A}]\,P[\bar{B}] \;-\; P[\bar{A}\bar{B}] \;=\; P[\bar{A}]\,P[\bar{B}] \;\geqslant\; 0.$$

To verify that (ii) $\Rightarrow$ (i), observe that if $C = \{\omega^*\}$ and $\omega^* \notin A$ in (17), then $SBC \neq \varnothing$ implies $SB \nsubseteq A$, and the left-hand side of (17) is again false, making the implication true.

Now assume condition (4) to be false, that is, $AB \neq \varnothing$ and $A \cup B \neq \Omega$. Pick $\omega_1 \in AB$ and $\omega_2 \in \bar{A}\bar{B}$; if $\omega^* \in A$, choose $\omega_1 = \omega^*$. Consider the following possibilistic and probabilistic knowledge worlds:

—If $C = \Omega$, consider the worlds (ω_1, S) and (ω_1, P) where $S = \{\omega_1, \omega_2\}$ and $P(\omega_1) = P(\omega_2) = 1/2$;

—If $C = \{\omega^*\}$, consider the worlds (ω^*, S) and (ω^*, P) where $S = \{\omega^*, \omega_1, \omega_2\}$ and P is uniform with support S. Note that $\omega^* \in A \;\Leftrightarrow\; |S| = 2$.

When $C = \{\omega^*\}$, in the possibilistic case we also assume that $\omega^* \in A$ (i.e., (ii) is false). Let us show that, for these worlds, Definitions 3.1 and 3.3 are both violated; that is:

$$SB \subseteq A \;\;\&\;\; S \nsubseteq A, \qquad P[A \mid B] \;>\; P[A].$$

The possibilistic part is obvious, since $SB = \{\omega_1\}$ and $S = \{\omega_1, \omega_2\}$. For the probabilistic part, if $|S| = 2$, then $P[A] = 1/2$ and $P[A \mid B] = 1$; if $|S| = 3$ and $\omega^* \notin A$, then $P(\omega_1) = P(\omega_2) = P(\omega^*) = 1/3$, and we have:

$$\begin{aligned} P[A \mid B] \;=\; P[AB] \,/\, P[B] \;&=\; P(\omega_1) \,/\, \big(P(\omega_1) + P(\omega^*)\big) \\ &=\; 1/2 \;>\; 1/3 \;=\; P(\omega_1) \;=\; P[A]. \quad \square \end{aligned}$$

Remark 3.13. In the auditing practice, the interesting case is $\omega^* \in A \cap B$, that is, when the protected and the disclosed properties are both true. In this case, unconditional privacy can be tested simply by checking whether $A \cup B = \Omega$, that is, whether "A or B" is always true.

4. *Possibilistic Case*

In this section, we shall focus exclusively on the possibilistic case; thus, the auditor's assumption about the user's knowledge shall be represented by $K \subseteq \Omega_{\text{poss}}$. While the probabilistic case is perhaps more interesting from the privacy perspective, the possibilistic case is simpler and provides intuition that sometimes extends to the probabilistic case. In fact, the possibilistic case is simple enough that useful statements can be proven in general, for arbitrary auditor's second-level knowledge sets $K \subseteq \Omega_{\text{poss}}$, or for a wide class of these sets.

Proposition 4.1 below gives a necessary and sufficient condition for K-preserving sets B to satisfy the privacy predicate $\text{Safe}_K(A, B)$, for a given and fixed set A. It associates every world $\omega \in A$ with a "safety margin" $\beta(\omega) \subseteq \Omega - A$ which depends only on ω, A and K. Given B, the condition verifies whether every $\omega \in A$ occurs in B together with its "safety margin," or does not occur in B at all. The "safety margin" ensures that this ω will not reveal A to the agent, no matter what prior knowledge $S \in \pi_2(K)$ the agent might have. (Recall that by π_i we denote the projection operation.)

PROPOSITION 4.1. *Let $K \subseteq \Omega_{\text{poss}}$ be an arbitrary second-level knowledge set, and assume $A \subseteq \Omega$. There exists a function $\beta\colon A \to \mathcal{P}(\Omega - A)$ such that $\forall B \subseteq \Omega$*

$$(\forall\, \omega \in AB :\ \beta(\omega) \subseteq B) \ \Rightarrow\ \text{Safe}_K(A, B), \tag{18}$$

and if B is K-preserving, then the converse holds:

$$\text{Safe}_K(A, B) \ \Rightarrow\ (\forall\, \omega \in AB :\ \beta(\omega) \subseteq B). \tag{19}$$

PROOF. For each $\omega \in A$, define $\beta(\omega)$ to be the $\bar{A}$-portion of the most informative K-preserving and K-privacy preserving disclosure $B_0(\omega)$ that is true at ω:

$$\beta(\omega) := B_0(\omega) - A, \text{ where } B_0(\omega) = \bigcap \left\{ B' \;\middle|\; \begin{array}{c} \omega \in B',\ \text{Safe}_K(A, B') \\ B' \text{ is } K\text{-preserving} \end{array} \right\}. \tag{20}$$

As an intersection of K-preserving sets B' that satisfy $\text{Safe}_K(A, B')$, by Proposition 3.10 the set $B_0(\omega)$ itself is K-preserving and satisfies $\text{Safe}_K(A, B_0(\omega))$.

To prove (18), let us assume $\forall\, \omega \in AB\colon \beta(\omega) \subseteq B$ and verify $\text{Safe}_K(A, B)$. Following Definition 3.1, we take some $(\omega, S) \in K$ such that $\omega \in B$ and $S \cap B \subseteq A$, and show that $S \subseteq A$. Since $\omega \in SB \subseteq A$, we have $\omega \in AB$, implying $\beta(\omega) \subseteq B$ by our assumption. We substitute $\beta(\omega) := B_0(\omega) - A$ and get $B_0(\omega) - A \subseteq B - A$, which in turn implies

$$S \cap B_0(\omega) - A \ \subseteq\ S \cap B - A \ =\ \varnothing,$$

that is, $S \cap B_0(\omega) \subseteq A$. By (20), we also have $\omega \in B_0(\omega)$. By the privacy definition for $B_0(\omega)$ we obtain $S \subseteq A$.

To prove (19), assume that B is K-preserving and satisfies $\text{Safe}_K(A, B)$; take an arbitrary $\omega \in AB$. Then, B is one of the sets intersected to define B_0 in (20), which gives us $\beta(\omega) \subseteq B_0 \subseteq B$. □

Remark 4.2. In the converse implication (19) of Proposition 4.1, we cannot drop the condition of B being K-preserving. Indeed, for all fixed A and β, the property $Q(B)$ defined as "$\forall\, \omega \in AB : \beta(\omega) \subseteq B$" is preserved under intersection: $Q(B_1)\ \&\ Q(B_2) \ \Rightarrow\ Q(B_1 \cap B_2)$. But $\text{Safe}_K(A, B)$, in general, is not preserved under intersection. For a simple example, let $\Omega = \{1, 2, 3\}$, $K = \Omega \otimes \{\Omega\}$, and $A = \{3\}$. Then both $B_1 = \{1, 3\}$ and $B_2 = \{2, 3\}$ protect the K-privacy of A, yet $B_1 \cap B_2 = \{3\}$ does not. However, see Theorem 4.14 for more on this subject.

The characterization in Proposition 4.1 could be quite useful for auditing a lot of properties $B_1, B_2, \ldots, B_N$ disclosed over a period of time, using the same audit query A. Given A, the auditor would compute the mapping β once, and use it to test every B_i. This comment applies to Section 4.1 as well.

4.1. INTERSECTION-CLOSED KNOWLEDGE

Motivation. When two or more possibilistic agents collude, that is, join forces in attacking protected information, their knowledge sets intersect: they jointly consider a world possible if and only if none of them has ruled it out. Therefore, if the auditor wants to account for potential collusions, she must consider knowledge world $(\omega, S_1 \cap S_2)$ possible whenever she considers both (ω, S_1) and (ω, S_2) possible. This motivates the following definition:

Definition 4.3. A second-level knowledge set $K \subseteq \Omega_{\text{poss}}$ is *intersection-closed*, or $\cap$-*closed* for short, when $\forall(\omega, S_1) \in K$ and $\forall(\omega, S_2) \in K$ we have $(\omega,\ S_1 \cap S_2) \in K$. Note that we intersect the user's knowledge sets (ω, S_1) and (ω, S_2) only when they are paired with the same world ω.

One way to obtain a second-level knowledge set $K \subseteq \Omega_{\text{poss}}$ that is $\cap$-closed is by taking an $\cap$-closed family Σ of subsets of Ω (such that $\forall S_1, S_2 \in \Sigma$: $S_1 \cap S_2 \in \Sigma$) and computing the product $K = C \otimes \Sigma$ with some knowledge set C.

Intervals. When the auditor's knowledge is $\cap$-closed, the notion of an "interval" between two worlds becomes central in characterizing the privacy relation:

Definition 4.4. Let $K \subseteq \Omega_{\text{poss}}$ be $\cap$-closed, and let $\omega_1, \omega_2 \in \Omega$ be two worlds such that

$$\omega_1 \in \pi_1(K), \quad \omega_2 \in \bigcup \{S \mid (\omega_1, S) \in K\}. \tag{21}$$

The K-*interval* from ω_1 to ω_2, denoted by $I_K(\omega_1, \omega_2)$, is the smallest set S such that $(\omega_1, S) \in K$ and $\omega_2 \in S$, or equivalently:

$$I_K(\omega_1, \omega_2) \ := \ \bigcap \{S \mid (\omega_1, S) \in K,\ \omega_2 \in S\}.$$

If the worlds ω_1, ω_2 do not satisfy conditions (21), we shall say that interval $I_K(\omega_1, \omega_2)$ does not exist.

Intuitively, $I_K(\omega_1, \omega_2)$ represents the "most knowledgeable" user who has not ruled out ω_2 when the actual world is ω_1. The following proposition shows that we need to know only the intervals in order to check whether or not $\text{Safe}_K(A, B)$ holds:

PROPOSITION 4.5. *For an $\cap$-closed set $K \subseteq \Omega_{\text{poss}}$ and for all $A, B \subseteq \Omega$, we have* $\text{Safe}_K(A, B)$ *if and only if*

$$\forall\, I_K(\omega_1, \omega_2): \quad \omega_1 \in AB \ \ \&\ \ \omega_2 \notin A \quad \Rightarrow \quad I_K(\omega_1, \omega_2) \cap (B - A) \ \neq \ \varnothing. \tag{22}$$

PROOF

(*if*) Assume (22) and let us prove $\text{Safe}_K(A, B)$. By Definition 3.1, we want to show

$$\forall(\omega, S) \in K\text{:}\ (\omega \in B \ \ \&\ \ S \cap B \subseteq A) \ \ \Rightarrow \ \ S \subseteq A. \tag{23}$$

Suppose that (23) is violated for $(\omega_1, S_1) \in K$; we have $\omega_1 \in AB$ and $\exists\, \omega_2 \in S_1 - A$. Interval $I_K(\omega_1, \omega_2) \subseteq S_1$ is well defined and satisfies the left-hand side of implication (22); hence, it satisfies the right-hand side too:

$$I_K(\omega_1, \omega_2) \cap (B - A) \ \neq \ \varnothing, \ \text{ which implies } \ S_1 \cap (B - A) \ \neq \ \varnothing.$$

But then (ω_1, S_1) does not violate (23) because the left-hand side of the implication (namely, $S_1 \cap B \subseteq A$) is false. Contradiction.

(*only if*) Assume $\text{Safe}_K(A, B)$, that is, (23), and let us prove (22). Take an arbitrary interval $S = I_K(\omega_1, \omega_2)$ such that $\omega_1 \in AB$ and $\omega_2 \notin A$, and consider a knowledge world $(\omega_1, S) \in K$. Since $\omega_2 \notin A$, we have $S \nsubseteq A$; to keep (23) true, we must also have $S \cap B \nsubseteq A$. This is the same as the right-hand side of (22). □

Remark 4.6. As implied by Proposition 4.5, there is no need to store the entire $\cap$-closed second-level knowledge set K (which could require $|\Omega| \cdot 2^{|\Omega|}$ bits of data) in

order to test the possibilistic privacy. It is sufficient to store one set $I_K(\omega_1, \omega_2) \subseteq \Omega$, or the fact of its non-existence, for each pair $(\omega_1, \omega_2) \in \Omega \times \Omega$, that is, at most $|\Omega|^3$ bits of data.

Minimal Intervals. In fact, in Proposition 4.5, we do not even have to check all intervals; it is enough to consider just "minimal" intervals defined as follows:

Definition 4.7. For an $\cap$-closed second-level knowledge set $K \subseteq \Omega_{\text{poss}}$, for a world $\omega_1 \in \Omega$ and for a set $X \subseteq \Omega$ not containing ω_1, an interval $I_K(\omega_1, \omega_2)$ is called a *minimal* K-interval *from* ω_1 *to* X when $\omega_2 \in X$ and

$$\forall\, \omega_2' \in X \cap I_K(\omega_1, \omega_2): \quad I_K(\omega_1, \omega_2') = I_K(\omega_1, \omega_2).$$

PROPOSITION 4.8. *For an $\cap$-closed set $K \subseteq \Omega_{\text{poss}}$ and for $\forall\, A, B \subseteq \Omega$, we have* $\text{Safe}_K(A, B)$ *if and only if the formula (22) holds over all intervals $I_K(\omega_1, \omega_2)$ that are minimal from a world $\omega_1 \in AB$ to the set $\Omega - A$.*

PROOF. We want to prove that, if (22) holds for all minimal intervals, then (22) holds for all intervals. It is sufficient to take an arbitrary interval $I_K(\omega_1, \omega_2)$ that satisfies $\omega_1 \in AB$ and $\omega_2 \in \bar{A}$, and show that it contains a minimal interval from ω_1 to $\bar{A}$. To find the minimal interval, start by setting $\omega_2^1 = \omega_2$, and continue to iteratively select ω_2^{n+1} given ω_2^n so that

$$\omega_2^{n+1} \in \bar{A} \cap I_K\big(\omega_1, \omega_2^n\big), \quad I_K\big(\omega_1, \omega_2^{n+1}\big) \subsetneq I_K\big(\omega_1, \omega_2^n\big)$$

until it is no longer possible, that is, until $I_K(\omega_1, \omega_2^n)$ is minimal. □

Example 4.9. Let Ω be an area of the plane that is bounded by a rectangle and discretized into pixels to ensure finiteness (the area within the 14×7 rectangle in Figure 1). Let the worlds be the pixels. Consider an auditor who does not know the actual database ω^* and who assumes that the user's prior knowledge set $S \in \Sigma$ is an integer rectangle, that is, a rectangle whose four corners have integer coordinates (corresponding to the vertical and horizontal lines in the picture). The family Σ of integer rectangles, and hence the auditor's second-level knowledge set $K = \Omega \otimes \Sigma$, are $\cap$-closed.

Given $\omega_1, \omega_2 \in \Omega$, the interval $I_K(\omega_1, \omega_2)$ is the smallest integer rectangle that contains both ω_1 and ω_2. For ω_1 and ω_2 in Figure 1, the interval $I_K(\omega_1, \omega_2)$ is the light-grey rectangle from point (1, 1) to point (4, 4); for ω_1 and ω_2', the interval $I_K(\omega_1, \omega_2')$ is the rectangle from point (1, 1) to point (9, 3).

The interval $I_K(\omega_1, \omega_2)$ shown on the picture is one of the three minimal intervals from ω_1 to set $\bar{A}$ (the area bounded by the ellipse). The other two minimal intervals are the rectangles (1, 1)−(5, 3) and (1, 1)−(6, 2). Every knowledge set S that the auditor considers possible in the case of $\omega^* = \omega_1$, that is, every S such that $(\omega_1, S) \in K$, must contain at least one of these three minimal intervals, unless $S \subseteq A$. For example, $S = I_K(\omega_1, \omega_2')$ in Figure 1 contains two minimal intervals (1, 1)−(5, 3) and (1, 1)−(6, 2). Thus, when looking for privacy violations, rather than going through all possible pairs $(\omega_1, S) \in K$ such that $\omega_1 \in B$ & $S \not\subseteq A$ and checking if $S \cap B \subseteq A$, the auditor has to go only through those (ω_1, S) that define minimal intervals to $\bar{A}$, a case of using Proposition 4.8.

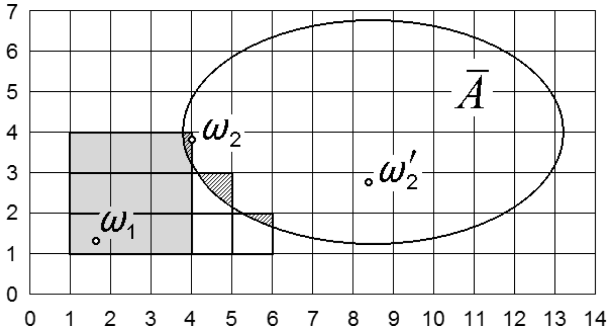

FIG. 1. An example of an $\cap$-closed $K \subseteq \Omega_{\text{poss}}$ where the worlds are the pixels inside the 14×7 rectangle (such as ω_1, ω_2 and ω_2'), and the permitted user's knowledge sets are the integer sub-rectangles (rectangles composed of whole squares). Set $\bar{A}$ is the complement of the privacy-sensitive knowledge set. See Example 4.9 for details.

Interval-Induced Partitions of $\bar{A}$. Let us have a closer look at the minimal K-intervals from a given world $\omega_1 \in A$ to the set $\bar{A} = \Omega - A$. For every $\omega_2 \in \bar{A}$, the interval $I_K(\omega_1, \omega_2)$, if it exists, is either minimal or not; if it is not minimal, then ω_2 cannot belong to any minimal interval from ω_1 to $\bar{A}$. Now, take some pair $\omega_2, \omega_2' \in \bar{A}$ such that both $I_K(\omega_1, \omega_2)$ and $I_K(\omega_1, \omega_2')$ are minimal. There are two possible situations:

(1) $I_K(\omega_1, \omega_2) = I_K(\omega_1, \omega_2')$, or
(2) $I_K(\omega_1, \omega_2) \cap I_K(\omega_1, \omega_2') \cap \bar{A} = \varnothing$.

Indeed, if $\exists\, \omega_2'' \in I_K(\omega_1, \omega_2) \cap I_K(\omega_1, \omega_2') \cap \bar{A}$, then by Definition 4.7 the interval $I_K(\omega_1, \omega_2'')$ equals both of the minimal intervals, making them equal. We have thus shown the following

PROPOSITION 4.10. *Given an $\cap$-closed set $K \subseteq \Omega_{\text{poss}}$, a set $A \subseteq \Omega$, and a world $\omega_1 \in A$, the minimal K-intervals from ω_1 to $\bar{A}$ partition set $\bar{A}$ into disjoint equivalence classes*

$$\bar{A} = D_1 \cup D_2 \cup \cdots \cup D_m \cup D'$$

where two worlds $\omega_2, \omega_2' \in \bar{A}$ belong to the same class D_i when they both belong to the same minimal interval, or (class D') when they both do not belong to any minimal interval.

Definition 4.11. In the assumptions and in the notation of Proposition 4.10, denote

$$\Delta_K(\bar{A}, \omega_1) := \{D_1, D_2, \ldots, D_m\}.$$

In other words, $\Delta_K(\bar{A}, \omega_1)$ is the disjoint collection of all sets formed by intersecting $\bar{A}$ with the minimal intervals from ω_1 to $\bar{A}$.

COROLLARY 4.12. *Given an $\cap$-closed set $K \subseteq \Omega_{\text{poss}}$, for all $A, B \subseteq \Omega$, we have* $\text{Safe}_K(A, B)$ *if and only if*

$$\forall\, \omega_1 \in AB, \quad \forall\, D_i \in \Delta_K(\bar{A}, \omega_1): \quad B \cap D_i \neq \varnothing. \tag{24}$$

PROOF. By Proposition 4.8, $\text{Safe}_K(A, B)$ holds if and only if for $\forall\, \omega_1 \in AB$ and for all intervals $I_K(\omega_1, \omega_2)$ that are minimal from ω_1 to $\bar{A}$ we have $I_K(\omega_1, \omega_2) \cap (B - A) \neq \varnothing$, or equivalently,

$$\forall\, \omega_1 \in AB,\ \forall\, I_K(\omega_1, \omega_2) \text{ minimal from } \omega_1 \text{ to } \bar{A}:\ \ B \cap (I_K(\omega_1, \omega_2) \cap \bar{A}) \neq \varnothing.$$

By Proposition 4.10, for every minimal $I_K(\omega_1, \omega_2)$ from ω_1 to $\bar{A}$, the intersection $I_K(\omega_1, \omega_2) \cap \bar{A}$ belongs to $\Delta_K(\bar{A}, \omega_1)$. Moreover, $\Delta_K(\bar{A}, \omega_1)$ contains all such intersections for the given $\bar{A}$ and ω_1, and contains nothing else. Replacing the quantifier over $I_K(\omega_1, \omega_2)$ with the quantifier over $D_i \in \Delta_K(\bar{A}, \omega_1)$ gives us (24). □

As Figure 1 illustrates for Example 4.9, the three minimal intervals from ω_1 to $\bar{A}$ formed by integer rectangles $(1, 1)-(4, 4)$, $(1, 1)-(5, 3)$ and $(1, 1)-(6, 2)$ are disjoint inside $\bar{A}$. Their intersections with $\bar{A}$, shown hatched in Figure 1, constitute the collection $\Delta_K(\bar{A}, \omega_1)$. A disclosed set B is private, assuming $\omega^* = \omega_1$, if and only if B intersects each of these three intervals inside $\bar{A}$.

The Case of All-Singleton $\boldsymbol{\Delta_K}$'s. If set K satisfies the property defined next,[5] privacy testing is simplified still further:

Definition 4.13. An $\cap$-closed set $K \subseteq \Omega_{\text{poss}}$ has *tight intervals* when for every K-interval $I_K(\omega_1, \omega_2)$ such that $\omega_1 \neq \omega_2$ we have

$$\forall\, \omega_2' \in I_K(\omega_1, \omega_2) - \{\omega_1, \omega_2\}:\quad I_K(\omega_1, \omega_2') \subsetneq I_K(\omega_1, \omega_2).$$

Informally, an interval from ω_1 to ω_2 is "tight" when for every point ω_2' in its "interior" the interval from ω_1 to ω_2' is strictly smaller (and hence no longer contains ω_2).

When K has tight intervals, every minimal interval $I_K(\omega_1, \omega_2)$ from $\omega_1 \in A$ to $\bar{A}$ has *exactly one* of its elements in $\bar{A}$, namely ω_2: $\bar{A} \cap I_K(\omega_1, \omega_2) = \{\omega_2\}$. Indeed, if $\bar{A} \cap I_K(\omega_1, \omega_2)$ contains another point $\omega_2' \neq \omega_2$, then $\omega_1 \notin \{\omega_2, \omega_2'\}$ since $\omega_1 \in A$, and by Definition 4.13 we get $I_K(\omega_1, \omega_2') \subsetneq I_K(\omega_1, \omega_2)$, that is, interval $I_K(\omega_1, \omega_2)$ is not minimal. Thus, for K that has tight intervals, all equivalence classes D_i in $\Delta_K(\bar{A}, \omega_1)$ are singletons, and Corollary 4.12 gives us the following characterization theorem (cf. Proposition 4.1):

THEOREM 4.14. *Let $K \subseteq \Omega_{\text{poss}}$ be an $\cap$-closed second-level knowledge set. The following three conditions are equivalent:*

(1) *K has tight intervals;*

(2) $\forall\, A \subseteq \Omega\ \ \exists\, \beta\colon A \to \mathcal{P}(\Omega - A)$ *such that* $\forall\, B \subseteq \Omega$*:*

$$\text{Safe}_K(A, B) \ \Leftrightarrow\ (\forall\, \omega \in AB:\ \beta(\omega) \subseteq B);$$

(3) $\forall\, A, B, B' \subseteq \Omega:\ \text{Safe}_K(A, B)\ \&\ \text{Safe}_K(A, B') \Rightarrow \text{Safe}_K(A, B \cap B')$*, that is, the privacy of individual disclosures always implies their joint privacy.*

[5] This definition slightly differs from the one given in the conference version: [Evfimievski et al. 2008].

PROOF

(1 $\Rightarrow$ 2): Let K have tight intervals, and assume $A \subseteq \Omega$. Define the function $\beta : A \to \mathcal{P}(\Omega - A)$ as given by

$$\forall\, \omega_1 \in A : \quad \beta(\omega_1) \;:=\; \bigcup \Delta_K(\bar{A}, \omega_1).$$

As we explained above, all D_i in the $\Delta_K(\bar{A}, \omega_1)$ of Corollary 4.12 are singletons, therefore $B \cap D_i \neq \varnothing$ is equivalent to $D_i \subseteq B$, and in (24)

$$\big(\forall\, D_i \in \Delta_K(\bar{A}, \omega_1) :\; B \cap D_i \neq \varnothing\big) \quad \Leftrightarrow \quad \bigcup \Delta_K(\bar{A}, \omega_1) \;\subseteq\; B.$$

(2 $\Rightarrow$ 3): If property "$\forall\, \omega \in AB :\; \beta(\omega) \subseteq B$" is satisfied for B and B', then it is also satisfied for $B \cap B'$. Indeed, take an arbitrary $\omega \in A \cap B \cap B'$, then $\omega \in AB$ implies $\beta(\omega) \subseteq B$ and $\omega \in AB'$ implies $\beta(\omega) \subseteq B'$; therefore, $\beta(\omega) \subseteq B \cap B'$. By Item 2, the property is equivalent to $\mathrm{Safe}_K(A, B)$.

(3 $\Rightarrow$ 1): We shall prove ($\neg 1 \Rightarrow \neg 3$) by assuming that K does not satisfy the tight intervals property (Definition 4.13) and constructing sets $A, B, B' \subseteq \Omega$ that violate Item 3. Let $I_K(\omega_1, \omega_2)$ be a "nontight" interval; that is, $\omega_1 \neq \omega_2$ and

$$\exists\, \omega_2' \;\in\; I_K(\omega_1, \omega_2) - \{\omega_1, \omega_2\} : \quad I_K(\omega_1, \omega_2') = I_K(\omega_1, \omega_2).$$

Notice that the three worlds ω_1, ω_2, and ω_2' are all different. Choose the sets as follows: $A = \Omega - \{\omega_2, \omega_2'\}$, $B = \{\omega_1, \omega_2\}$, and $B' = \{\omega_1, \omega_2'\}$. Then, we have:

— $AB = AB' = ABB' = \{\omega_1\}$;

— $I \;:=\; I_K(\omega_1, \omega_2) \;=\; I_K(\omega_1, \omega_2')$ is the only minimal interval from ω_1 to $\{\omega_2, \omega_2'\} = \bar{A}$;

— $I \cap (B - A) = \{\omega_2\}$, $I \cap (B' - A) = \{\omega_2'\}$, and $I \cap (BB' - A) = \varnothing$.

By Proposition 4.8, we have $\mathrm{Safe}_K(A, B)$, $\mathrm{Safe}_K(A, B')$, but not $\mathrm{Safe}_K(A, BB')$. □

In practice, Condition 3 in Theorem 4.14 is very desirable: it allows the auditor to verify the safety of a sequence $B_1, \ldots, B_k$ of disclosed queries by testing each query individually, even though the auditor's prior assumptions K about the user's knowledge no longer hold after some or all of the disclosures. For example, if the disclosure of subsequence $B_1, \ldots, B_{k-1}$ protects the privacy of a certain database property A, but the disclosure of the entire sequence violates it, then Condition 3 for K implies $\neg\, \mathrm{Safe}_K(A, B_k)$; the same is true for any other subsequence of the disclosed queries.

Several important examples of second-level knowledge sets that have tight intervals are discussed in Section 5.1. See Remark 4.2 for a counterexample where an $\cap$-closed K does not have tight intervals.

Remark 4.15. When the auditor knows that the actual database is precisely ω^*, her second-level knowledge set K contains only knowledge worlds of the form (ω^*, S). Then all collections $\Delta_K(\bar{A}, \omega_1)$ and sets $\beta(\omega_1)$ are empty for all $\omega_1 \neq \omega^*$ because there exist no intervals $I_K(\omega_1, \omega_2)$, and we have to check only the case of $\omega_1 = \omega^*$ in the above privacy tests.

5. *Modularity Assumptions for Probabilistic Knowledge*

In the previous section, we clarified some general properties of possibilistic knowledge; now we turn to the more complex probabilistic case. Rather than studying arbitrary probabilistic knowledge families, here we shall focus on a few specific, yet important, families of distributions. We shall also see some concrete examples of possibilistic knowledge families induced by the probabilistic ones. Later, in Section 6, we present more sophisticated approaches that extend beyond these families.

From now on, we assume that $\Omega = \{0,1\}^n$ for some fixed n. Let $\omega_1 \wedge \omega_2$ ($\omega_1 \vee \omega_2$, $\omega_1 \oplus \omega_2$) be the bit-wise "AND" ("OR", "XOR"), and define the partial order $\omega_1 \leqslant \omega_2$ to mean "$\forall\, i = 1, \ldots, n$: $\omega_1[i] = 1 \;\Rightarrow\; \omega_2[i] = 1$."

Definition 5.1. A probability distribution P over Ω is called *log-supermodular* (*log-submodular*)[6] when the following holds:

$$\forall\, \omega_1, \omega_2 \in \Omega:\;\; P(\omega_1)\, P(\omega_2) \;\leqslant (\geqslant)\;\; P(\omega_1 \wedge \omega_2)\, P(\omega_1 \vee \omega_2).$$

The family of all log-supermodular distributions shall be denoted by $\Pi^+_{\mathfrak{m}}$, the family of all log-submodular distributions by $\Pi^-_{\mathfrak{m}}$.

A distribution P is called a *product distribution* if it makes every coordinate independent. Every product distribution corresponds to a vector $(p_1, \ldots, p_n)$ of Bernoulli probabilities, each $p_i \in [0,1]$, such that

$$\forall\, \omega \in \{0,1\}^n:\;\; P(\omega) \;=\; \textstyle\prod_{i=1}^n \;\; p_i^{\omega[i]} \cdot (1 - p_i)^{1-\omega[i]}. \tag{25}$$

The family of all product distributions shall be denoted by $\Pi^0_{\mathfrak{m}}$.

PROPOSITION 5.2. *We have* $\Pi^0_{\mathfrak{m}} = \Pi^-_{\mathfrak{m}} \cap \Pi^+_{\mathfrak{m}}$. *Equivalently, P is a product distribution if and only if*

$$\forall\, \omega_1, \omega_2 \in \Omega:\;\; P(\omega_1)\, P(\omega_2) \;=\; P(\omega_1 \wedge \omega_2)\, P(\omega_1 \vee \omega_2). \tag{26}$$

PROOF. This is a minor variation of a statement proven in Lovász [1983]; we include the proof below for the sake of completeness.

For every $i = 1, \ldots, n$ the bit pair $\omega_1[i], \omega_2[i]$ contains the same number of 0's and 1's as the bit pair $(\omega_1 \wedge \omega_2)[i]$, $(\omega_1 \vee \omega_2)[i]$. Therefore, if P is a product distribution, then $\forall\, i = 1, \ldots, n$ the terms p_i and $1 - p_i$ appear the same number of times on the left-hand side and on the right-hand side of (26), making the sides equal.

Conversely, (26) implies $\forall\, \omega, \omega' \in \Omega$, $\forall\, i = 1, \ldots, n$:

$$P\big(\omega|_{\omega[i]\leftarrow 0}\big) \cdot P\big(\omega'|_{\omega'[i]\leftarrow 1}\big) \;=\; P\big(\omega|_{\omega[i]\leftarrow 1}\big) \cdot P\big(\omega'|_{\omega'[i]\leftarrow 0}\big), \tag{27}$$

where "$\omega[i] \leftarrow b$" means "set the ith bit in ω to b." As a probability function, P must sum up to 1, hence P must be nonzero at some $\omega' \in \Omega$. Take an arbitrary $i = 1, \ldots, n$ and assume that $\omega'[i] = 0$; then we can rewrite (27) as

$$\forall\, \omega \in \Omega:\;\; P\big(\omega|_{\omega[i]\leftarrow 1}\big) \;=\; c_i \cdot P\big(\omega|_{\omega[i]\leftarrow 0}\big), \quad c_i \;=\; P\big(\omega'|_{\omega'[i]\leftarrow 1}\big) \,\big/\, P\big(\omega'|_{\omega'[i]\leftarrow 0}\big).$$

Set $p_i = c_i/(1 + c_i)$. If instead we have $\omega'[i] = 1$, then rewrite (27) as

$$\forall\, \omega \in \Omega:\;\; P\big(\omega|_{\omega[i]\leftarrow 0}\big) \;=\; c'_i \cdot P\big(\omega|_{\omega[i]\leftarrow 1}\big), \quad c'_i \;=\; P\big(\omega'|_{\omega'[i]\leftarrow 0}\big) \,\big/\, P\big(\omega'|_{\omega'[i]\leftarrow 1}\big),$$

[6] The "log-" means that supermodularity is multiplicative, rather than additive. The subscript "$\mathfrak{m}$" in $\Pi^-_{\mathfrak{m}}$, $\Pi^+_{\mathfrak{m}}$ etc. means "modular."

and set $p_i = 1/(1 + c'_i)$. By induction on the Hamming distance of ω from ω', we can check that every $P(\omega)$ is proportional to the product distribution (25). Therefore, P is a product distribution. □

Supermodular and submodular functions occur often in mathematics and have been extensively studied [Fujishige 2005; Lovász 1983]. Our goal in considering these assumptions was to substantially relax bit-wise independence while staying away from the unconstrained case. Besides that, the log-supermodular assumption (as implied by Theorem 5.10 in Section 5.2) describes situations where no negative correlations are permitted across individual database records—something we might expect from knowledge about, say, HIV incidence among humans. The following example provides a case in point:

Example 5.3. Let us consider a probability distribution $P : \Omega \to \mathbb{R}_+$ that has the form

$$P(\omega) \;=\; C\,\exp\left(\sum_{i=1}^{n} a_i\,\omega[i] \;+\sum_{1\leqslant i<j\leqslant n} b_{i,j}\,\omega[i]\,\omega[j]\right), \quad \text{where } \forall\, i,\, j:\; b_{i,j} \geqslant 0. \tag{28}$$

The log-linear expression in (28) naturally arises when P is the maximum entropy distribution with equality constraints on single-bit expectations and two-bit covariances [Cover and Thomas 2006]. It is used extensively in machine learning, for example in the definition of the Boltzmann machine [Ackley et al. 1985], but without our requirement that all $b_{i,j}$ be nonnegative.

It is easy to see that a distribution of the form (28) is always log-supermodular. Indeed, since $C\exp(x)\cdot C\exp(y) = C^2\exp(x+y)$, for all ω_1 and ω_2 in Ω we have:

$$\begin{aligned} P(\omega_1)\,P(\omega_2) \;=\; C^2\exp\Bigg(&\sum_{i=1}^{n} a_i\big(\omega_1[i]+\omega_2[i]\big) \\ &+\sum_{1\leqslant i<j\leqslant n} b_{i,j}\big(\omega_1[i]\,\omega_1[j]+\omega_2[i]\,\omega_2[j]\big)\Bigg) \end{aligned}$$

$$P(\omega_1\wedge\omega_2)\,P(\omega_1\vee\omega_2) = C^2\exp\Bigg(\sum_{i=1}^{n} a_i\,\big((\omega_1\wedge\omega_2)[i]\;+\;(\omega_1\vee\omega_2)[i]\big)$$

$$+\sum_{1\leqslant i<j\leqslant n} b_{i,j}\,\big((\omega_1\wedge\omega_2)[i]\,(\omega_1\wedge\omega_2)[j]+(\omega_1\vee\omega_2)[i]\,(\omega_1\vee\omega_2)[j]\big)\Bigg),$$

where, because ω_1 and ω_2 are from $\Omega = \{0,1\}^n$, for all i and j we always have

$$\begin{aligned} \omega_1[i]+\omega_2[i] \;&=\; (\omega_1\wedge\omega_2)[i]+(\omega_1\vee\omega_2)[i] \\ \omega_1[i]\omega_1[j]\;+\;\omega_2[i]\omega_2[j] \;&\leqslant\; (\omega_1\wedge\omega_2)[i](\omega_1\wedge\omega_2)[j] \\ &\quad +(\omega_1\vee\omega_2)[i]\,(\omega_1\vee\omega_2)[j], \end{aligned}$$

which gives us $P(\omega_1)\,P(\omega_2) \leqslant P(\omega_1\wedge\omega_2)\,P(\omega_1\vee\omega_2)$, since all $b_{i,j}\geqslant 0$.

5.1. MODULARITY FOR SETS. Let us define three families of sets composed of the supports of all distributions in Π^-_{m}, Π^+_{m}, and Π^0_{m}:

$$\Sigma^-_{\mathrm{m}} = \mathrm{supp}\left(\Pi^-_{\mathrm{m}}\right), \quad \Sigma^+_{\mathrm{m}} = \mathrm{supp}\left(\Pi^+_{\mathrm{m}}\right), \quad \Sigma^0_{\mathrm{m}} = \mathrm{supp}\left(\Pi^0_{\mathrm{m}}\right);$$

here, as before, $\mathrm{supp}(\Pi)$ denotes $\{\mathrm{supp}(P) \mid P \in \Pi\}$. These families of sets have a simple characterization, given in the following definition and in Propositions 5.6 and 5.7, which we now derive.

Definition 5.4. A set $S \subseteq \Omega$ is an *up-set* (a *down-set*) when $\forall\, \omega_1 \in S, \forall\, \omega_2 \geqslant \omega_1$ ($\forall\, \omega_2 \leqslant \omega_1$) we have $\omega_2 \in S$. A nonempty intersection of an up-set and a down-set shall be called a *convex set*. A nonempty set $S \subseteq \Omega$ is a *sublattice* when

$$\forall\, \omega_1, \omega_2 \in S: \quad \omega_1 \wedge \omega_2 \in S \text{ and } \omega_1 \vee \omega_2 \in S. \tag{29}$$

A nonempty set $S \subseteq \Omega$ is a *product set* when

$$S = S_1 \times S_2 \times \cdots \times S_n, \quad S_i = \{0\} \text{ or } \{1\} \text{ or } \{0, 1\}.$$

Remark 5.5. An intersection of up-sets is an up-set, of down-sets is a down-set; set $S \subseteq \Omega$ is an up-set if and only if $\bar{A}$ is a down-set. A nonempty intersection of convex sets is a convex set, of sublattices is a sublattice, of product sets is a product set.

PROPOSITION 5.6

(*a*) *A nonempty set* $S \subseteq \Omega$ *is convex if and only if*

$$\forall\, \omega_1, \omega_2 \in S, \ \forall\, \omega \in \Omega: \quad \omega_1 \leqslant \omega \leqslant \omega_2 \ \Rightarrow\ \omega \in S. \tag{30}$$

(*b*) *A nonempty set* $S \subseteq \Omega$ *is a sublattice if and only if the property "*$\omega \in S$*" can be expressed as a conjunction of two-bit implications*[7] *of the form "*$\omega[i] \to \omega[j]$*" and one-bit lookups of the form "*$\omega[i] = 0$*" or "*$\omega[i] = 1$*."*

PROOF

(a) An intersection of an up-set U and a down-set D must satisfy (30) because $\omega_1 \in U$ implies $\omega \in U$ and $\omega_2 \in D$ implies $\omega \in D$. Conversely, every set that satisfies (30) can be represented as such an intersection $U \cap D$ as follows:

$$U = \{\omega \in \Omega \mid \exists\, \omega_1 \in S: \ \omega_1 \leqslant \omega\}, \quad D = \{\omega \in \Omega \mid \exists\, \omega_2 \in S: \ \omega \leqslant \omega_2\}.$$

(b) For the "if" direction, it is easy to see that sets $\{\omega \in \Omega \mid \omega[i] \to \omega[j]\}$ and $\{\omega \in \Omega \mid \omega[i] = b\}$ are sublattices, for all i and j; a conjunction of such implications and lookups gives an intersection of sublattices, which is also a sublattice (if nonempty). A straightforward proof for the "only if" direction by induction on n is a bit tedious, so we instead refer to Table 2 of Creignou et al. [2008]. The set of all sublattices over $\{0, 1\}^n$ is a special case of *co-clone*, the notion studied in that paper. Given a set $\mathcal{F}$ of Boolean functions, the co-clone $\mathrm{Inv}(\mathcal{F})$ is the collection of all subsets $S \subseteq \{0, 1\}^n$ (for some n) that satisfy

$$\forall f \in \mathcal{F}, \ m := \mathrm{arity}(f), \quad \forall\, \omega_1, \ldots, \omega_m \in S: \quad f(\omega_1, \ldots, \omega_m) \in S,$$

where

$$f(\omega_1, \ldots, \omega_m) := \omega \in \{0, 1\}^n \text{ such that}$$
$$\forall\, i = 1, \ldots, n: \ \omega[i] = f(\omega_1[i], \ldots, \omega_m[i]).$$

[7]An implication "$\omega[i] \to \omega[j]$" is the same as formula "$\neg\, \omega[i] \vee \omega[j]$".

Informally, $S \in \mathrm{Inv}(\mathcal{F})$ means S is "preserved" under all Boolean operations in $\mathcal{F}$ applied bit-wise to vectors in S. In particular [Böhler et al. 2003], co-clone $IM_2 = \mathrm{Inv}(\{\wedge, \vee\})$ gives the set of all sublattices, as defined by (29). Table 2 in Creignou et al. [2008] gives a "plain basis" for every Boolean co-clone, that is, a set of Boolean relations whose conjunctions generate precisely all subsets in the co-clone. In our special case, it shows that IM_2 is generated by two-bit implications and single-bit lookups. □

PROPOSITION 5.7. *The following equalities hold:*

— $\Sigma_{\mathrm{m}}^{-} = \{$*all convex sets over* $\Omega\}$;
— $\Sigma_{\mathrm{m}}^{+} = \{$*all sublattices over* $\Omega\}$;
— $\Sigma_{\mathrm{m}}^{0} = \{$*all product sets over* $\Omega\} = \Sigma_{\mathrm{m}}^{+} \cap \Sigma_{\mathrm{m}}^{-}$.

PROOF. By Proposition 5.6, a set $S \neq \varnothing$ is convex if and only if $\forall u, v \in S$: $u \leqslant \omega \leqslant v \Rightarrow \omega \in S$. We now show that this is equivalent to

$$\forall \omega_1, \omega_2 \in \Omega\colon \{\omega_1 \wedge \omega_2,\ \omega_1 \vee \omega_2\} \subseteq S \Rightarrow \{\omega_1, \omega_2\} \subseteq S. \tag{31}$$

Indeed, for a convex S the above implication holds because $\omega_1 \wedge \omega_2 \leqslant \omega_i \leqslant \omega_1 \vee \omega_2$ for $i = 1, 2$. Now let us assume (31), take some $u, v \in S$ and $u \leqslant \omega \leqslant v$, and show $\omega \in S$. Define $\omega' = \omega \oplus u \oplus v$, that is, we have $\omega'[i] = \omega[i]$ iff $u[i] = v[i]$. It is not hard to verify that $u = \omega \wedge \omega'$ and $v = \omega \vee \omega'$, so by (31) $u, v \in S$ implies $\omega, \omega' \in S$.

Given a nonempty set S, define a probability distribution P_S to be identical (uniform) on all $\omega \in S$ and zero everywhere else. For a convex S, distribution P_S is log-submodular due to (31): $\forall \omega_1, \omega_2 \in \Omega$,

$$P(\omega_1 \wedge \omega_2)P(\omega_1 \vee \omega_2) \neq 0 \Rightarrow P(\omega_1)P(\omega_2) = 1/|S|^2 = P(\omega_1 \wedge \omega_2)P(\omega_1 \vee \omega_2).$$

Since $S = \mathrm{supp}(P_S)$, we obtain $S \in \Sigma_{\mathrm{m}}^{-}$. Conversely, if P is log-submodular, then (31) must hold for $S = \mathrm{supp}(P)$ in order to satisfy Definition 5.1, proving the convexity of $\mathrm{supp}(P)$.

In the same way, given a sublattice S, the distribution P_S is log-supermodular due to (29) in Definition 5.4, and conversely, $\forall P \in \Pi_{\mathrm{m}}^{+}$ the set $\mathrm{supp}(P)$ has to be a sublattice in order to satisfy Definition 5.1. Lastly, for a product set S, the distribution P_S is a product distribution with vector $(p_1, \ldots, p_n)$ where all $p_i \in \{0, 1, 1/2\}$, and conversely, supports of product distributions must be product sets. Equality $\Sigma_{\mathrm{m}}^{0} = \Sigma_{\mathrm{m}}^{+} \cap \Sigma_{\mathrm{m}}^{-}$ for sets S is implied by $\Pi_{\mathrm{m}}^{0} = \Pi_{\mathrm{m}}^{+} \cap \Pi_{\mathrm{m}}^{-}$ for distributions P_S. □

Families Σ_{m}^{-}, Σ_{m}^{+} and Σ_{m}^{0} can also be viewed as possibilistic knowledge assumptions. For example, the family Σ_{m}^{-} of convex sets describes a user's possibilistic knowledge about the actual database ω^* learned by issuing a sequence of monotone Boolean queries[8] and receiving "yes" or "no" answers. Family Σ_{m}^{0} of product sets describes the possibilistic knowledge learned by asking, for a sequence of records, whether or not each given record belongs to the database. All three families are

[8] A monotone Boolean query is a mapping $Q : \Omega \to \{\text{"yes"}, \text{"no"}\}$ such that $\forall \omega_1, \omega_2 \in \Omega$ if $Q(\omega_1) =$ "yes" and $\omega_1 \leqslant \omega_2$ then $Q(\omega_2) =$ "yes".

$\cap$-closed, barring the empty intersections (see Remark 5.5). Therefore, for every set $C \neq \varnothing$ the second-level knowledge sets $C \otimes \Sigma_{\mathrm{m}}^{-}$, $C \otimes \Sigma_{\mathrm{m}}^{+}$ and $C \otimes \Sigma_{\mathrm{m}}^{0}$ are intersection closed, and Section 4.1 applies. Let us compute for them the intervals introduced in Definition 4.4:

PROPOSITION 5.8. *Assuming $\omega_1 \in C$, we have:*

$$\begin{aligned} I_{C\otimes\Sigma_{\mathrm{m}}^{-}}(\omega_1,\omega_2) &= \begin{cases} \{\omega \mid \omega_1 \leqslant \omega \leqslant \omega_2\}, & \text{if } \omega_1 \leqslant \omega_2, \\ \{\omega \mid \omega_2 \leqslant \omega \leqslant \omega_1\}, & \text{if } \omega_2 \leqslant \omega_1, \\ \{\omega_1,\ \omega_2\}, & \text{if } \omega_1 \not\leqslant \omega_2 \text{ and } \omega_1 \not\geqslant \omega_2; \end{cases} \\ I_{C\otimes\Sigma_{\mathrm{m}}^{+}}(\omega_1,\omega_2) &= \{\omega_1,\ \omega_2,\ \omega_1\wedge\omega_2,\ \omega_1\vee\omega_2\}; \\ I_{C\otimes\Sigma_{\mathrm{m}}^{0}}(\omega_1,\omega_2) &= \{\omega \mid \omega_1\wedge\omega_2 \leqslant \omega \leqslant \omega_1\vee\omega_2\}. \end{aligned} \tag{32}$$

All these second-level knowledge sets satisfy the "tight intervals" property (see Definition 4.13*), and therefore the items of Theorem* 4.14 *apply to them.*

PROOF. First, let us get convinced that the sets on the right-hand side of the above equalities (32) belong to their respective families Σ_{m}^{-}, Σ_{m}^{+} and Σ_{m}^{0}. Indeed, for all $\omega' \leqslant \omega''$ the set $\{\omega \mid \omega' \leqslant \omega \leqslant \omega''\}$ is convex as an intersection of an up-set and a down-set, and it is a sublattice too, because operations $\wedge$ and $\vee$ respect a common lower or upper bound; hence, it is a product set (Proposition 5.7). A set $\{\omega_1,\ \omega_2\}$ of two (or any number of) incomparable worlds is convex because it satisfies the implication in (30), while set $\{\omega_1,\ \omega_2,\ \omega_1\wedge\omega_2,\ \omega_1\vee\omega_2\}$ is the sublattice generated by ω_1 and ω_2.

Second, let us show that the sets on the right-hand side of (32) are subsets of all sets that contain ω_1 and ω_2 from their respective families Σ_{m}^{-}, Σ_{m}^{+} and Σ_{m}^{0}; this will prove that these sets satisfy Definition 4.4. If a convex set contains ω_1 and ω_2 where $\omega_1 \leqslant \omega_2$ or $\omega_2 \leqslant \omega_1$, then by (30) the set contains everything between ω_1 and ω_2. If a sublattice contains ω_1 and ω_2, then by definition it contains $\omega_1\wedge\omega_2$ and $\omega_1\vee\omega_2$. If a product set contains ω_1 and ω_2, then as a sublattice it contains $\omega_1\wedge\omega_2$ and $\omega_1\vee\omega_2$, and as a convex set it contains everything in between. This proves that the right-hand sides are indeed the intervals between ω_1 and ω_2.

Finally, let us show that these intervals are "tight" by verifying Definition 4.13. We consider each family in turn:

$K = C \otimes \Sigma_{\mathrm{m}}^{-}$ If ω_1 and ω_2 are comparable, say $\omega_1 \leqslant \omega_2$, and if we pick some world $\omega_2' \notin \{\omega_1, \omega_2\}$ from $I_K(\omega_1,\omega_2) = \{\omega \mid \omega_1 \leqslant \omega \leqslant \omega_2\}$ and construct $I_K(\omega_1,\omega_2') = \{\omega \mid \omega_1 \leqslant \omega \leqslant \omega_2'\}$, the new interval will not contain ω_2. If ω_1 and ω_2 are incomparable, the original interval is $\{\omega_1, \omega_2\}$ and there is nothing to pick as ω_2'.

$K = C \otimes \Sigma_{\mathrm{m}}^{+}$ If ω_1 and ω_2 are incomparable, the original interval contains four different worlds, and picking, say, $\omega_2' = \omega_1\wedge\omega_2$ reduces the interval to two worlds. If ω_1 and ω_2 are comparable, we start out with a two-world interval, so there is nothing to pick as ω_2'.

$K = C \otimes \Sigma_{\mathrm{m}}^{0}$ The original interval $I_K(\omega_1,\omega_2) = \{\omega \mid \omega_1\wedge\omega_2 \leqslant \omega \leqslant \omega_1\vee\omega_2\}$ can be equivalently written as

$$I_K(\omega_1,\omega_2) = \{\omega \mid \forall\, i = 1,\ldots,n:\ \omega_1[i]=\omega_2[i] \ \Rightarrow\ \omega[i]=\omega_1[i]=\omega_2[i]\}.$$

If we pick some $\omega_2' \neq \omega_2$ from this interval, the set of bit indices $\{i \mid \omega_1[i] = \omega_2'[i]\}$ will be a strict superset of $\{i \mid \omega_1[i] = \omega_2[i]\}$, and therefore ω_2 will not make it into the new interval $I_K(\omega_1, \omega_2')$. □

5.2. PRIVACY FOR LOG-SUPERMODULAR DISTRIBUTIONS. Let us come back to the probabilistic knowledge, specifically to the three families of distributions introduced by Definition 5.1: Π_m^+ (log-supermodular distributions), Π_m^- (log-submodular distributions), and Π_m^0 (product distributions). We shall be interested in necessary criteria and in sufficient criteria for testing privacy over these families. From here onwards, the probabilistic privacy will be understood in the sense of Definition 3.5.

One way to produce a necessary criterion for probabilistic privacy is by converting the family of probabilities into a possibilistic family of supports of these probabilities, as we discussed in Section 3.3. We can then consider the privacy test for this possibilistic family, and use the implication (15), repeated next:

$$\text{Safe}_{\Pi}(A, B) \;\Rightarrow\; \text{Safe}_{\Omega\otimes\Sigma}(A, B) \;\&\; \text{Safe}_{\Omega\otimes\Sigma}(\bar{A}, \bar{B}), \tag{33}$$

where $\Sigma = \text{supp}(\Pi)$. Let us instantiate this criterion for the family Π_m^+ of log-supermodular distributions:

PROPOSITION 5.9 (Π_m^+ SAFETY: NECESSARY CRITERION). *For all* $A, B \subseteq \Omega = \{0, 1\}^n$ *such that* $\text{Safe}_{\Pi_m^+}(A, B)$*, every pair of worlds* $\omega_1 \in AB$ *and* $\omega_2 \in \bar{A}\bar{B}$ *satisfies one of the following two conditions:*

— $\omega_1 \wedge \omega_2 \in A - B$ *and* $\omega_1 \vee \omega_2 \in B - A$*;*
— $\omega_1 \wedge \omega_2 \in B - A$ *and* $\omega_1 \vee \omega_2 \in A - B$.

PROOF. By definition and by Proposition 5.7, we have $\text{supp}(\Pi_m^+) = \Sigma_m^+$, the family of all sublattices. In order to apply (33), we need a test for the possibilistic privacy predicate $\text{Safe}_K(A, B)$, where $K = \Omega \otimes \Sigma_m^+$. Since K is $\cap$-closed, let us use the interval-based test given by Proposition 4.5: $\text{Safe}_K(A, B)$ if and only if

$$\forall\, \omega_1 \in AB,\;\; \forall\, \omega_2 \notin A:\;\; I_{\Omega\otimes\Sigma_m^+}(\omega_1, \omega_2) \cap (B - A) \neq \varnothing,$$

where we can restrict ω_2 to set $\bar{A}\bar{B}$, since for $\omega_2 \in B - A$ the formula is vacuously true. From Proposition 5.8, we know that

$$I_{\Omega\otimes\Sigma_m^+}(\omega_1, \omega_2) = \{\omega_1,\; \omega_2,\; \omega_1 \wedge \omega_2,\; \omega_1 \vee \omega_2\};$$

therefore, we have $\text{Safe}_K(A, B)$ if and only if

$$\forall\, \omega_1 \in AB,\;\; \forall\, \omega_2 \in \bar{A}\bar{B}:\;\; \{\omega_1 \wedge \omega_2,\; \omega_1 \vee \omega_2\} \cap (B - A) \neq \varnothing.$$

Analogously, we have $\text{Safe}_K(\bar{A}, \bar{B})$ if and only if

$$\forall\, \omega_1 \in AB,\;\; \forall\, \omega_2 \in \bar{A}\bar{B}:\;\; \{\omega_1 \wedge \omega_2,\; \omega_1 \vee \omega_2\} \cap (A - B) \neq \varnothing.$$

Substituting these tests into (33) for $\Pi = \Pi_m^+$ and $\Sigma = \Sigma_m^+$ completes the proof. □

It turns out that one can prove a sufficient criterion for Π_m^+-safety that has a form very similar to Proposition 5.9, although not quite the same. The sufficient criterion relies on the following well-known theorem introduced in Ahlswede and Daykin [1978]:

THEOREM 5.10 (FOUR FUNCTIONS THEOREM). *Let L be a finite distributive lattice,*[9] *and let $\alpha, \beta, \gamma, \delta : L \to \mathbb{R}_+$. For all subsets $A, B \subseteq L$ denote $f[A] = \sum_{a \in A} f(a)$, $A \vee B = \{a \vee b \,|\, a \in A,\ b \in B\}$, and $A \wedge B = \{a \wedge b \,|\, a \in A,\ b \in B\}$. Then, the inequality*

$$\alpha[A] \cdot \beta[B] \ \leqslant\ \gamma[A \vee B] \cdot \delta[A \wedge B]$$

holds for all subsets $A, B \subseteq L$ if and only if it holds for one-element subsets, that is, if and only if

$$\alpha(a) \cdot \beta(b) \ \leqslant\ \gamma(a \vee b) \cdot \delta(a \wedge b)$$

for all elements $a, b \in L$.

PROOF. See, for example, Bollobás [1986, Section 19]. □

PROPOSITION 5.11 (Π^+_{m} SAFETY: SUFFICIENT CRITERION). *For all $A, B \subseteq \Omega = \{0,1\}^n$, either one of the two conditions below is sufficient to establish* $\mathrm{Safe}_{\Pi^+_{\mathrm{m}}}(A, B)$:

— $AB \wedge \bar{A}\bar{B} \subseteq A - B$ *and* $AB \vee \bar{A}\bar{B} \subseteq B - A$;
— $AB \wedge \bar{A}\bar{B} \subseteq B - A$ *and* $AB \vee \bar{A}\bar{B} \subseteq A - B$.

PROOF. Let $P \in \Pi^+_{\mathrm{m}}$, set the four functions as $\alpha = \beta = \gamma = \delta = P$, and set the distributive lattice $L = \Omega = \{0,1\}^n$. The log-supermodularity definition and Theorem 5.10 imply $\forall\, A, B \subseteq \Omega$

$$\begin{aligned} P[AB] \cdot P[\bar{A}\bar{B}] \ &\leqslant\ P[AB \vee \bar{A}\bar{B}] \cdot P[AB \wedge \bar{A}\bar{B}] \\ &\leqslant\ P[A - B] \cdot P[B - A], \end{aligned}$$

where the last "$\leqslant$" is implied by either of the two conditions assumed in our proposition. It remains to recall that by Proposition 3.8

$$P[A]\,P[B] \ -\ P[AB] \ =\ P[A - B] \cdot P[B - A] \ -\ P[AB] \cdot P[\bar{A}\bar{B}],$$

and the definition of $\mathrm{Safe}_{\Pi}(A, B)$ given by (10). □

COROLLARY 5.12. *If A is an up-set and B is a down-set (or vice-versa), then* $\mathrm{Safe}_{\Pi^+_{\mathrm{m}}}(A, B)$.

PROOF. Let us show that, if A and $\bar{B}$ are both up-sets, then $A \vee \bar{B} = A\bar{B} = A - B$. Indeed, $\forall\, \omega \in A \vee \bar{B}$ we have $\omega = a \vee b'$ where $a \in A$ and $b' \in \bar{B}$, implying $a \leqslant \omega \in A$ and $b' \leqslant \omega \in \bar{B}$; on the other hand, $\forall\, \omega \in A \cap \bar{B}$ we have $\omega = \omega \vee \omega \in A \vee \bar{B}$. Analogously, since $\bar{A}$ and B are down-sets, also $B \wedge \bar{A} = B - A$. We have:

$$\begin{aligned} AB \subseteq A \ \&\ \bar{A}\bar{B} \subseteq \bar{B} \ &\Rightarrow\ AB \vee \bar{A}\bar{B} \ \subseteq\ A \vee \bar{B} \ =\ A - B; \\ AB \subseteq B \ \&\ \bar{A}\bar{B} \subseteq \bar{A} \ &\Rightarrow\ AB \wedge \bar{A}\bar{B} \ \subseteq\ B \wedge \bar{A} \ =\ B - A. \end{aligned}$$

The rest follows from Proposition 5.11. If it is B that is the up-set, and A is the down-set, just permute A and B everywhere in the proof. □

[9] A *lattice* L is a partially ordered set where every pair of elements $a, b \in L$ has the least upper bound $a \vee b$ and the greatest lower bound $a \wedge b$. A lattice is *distributive* when $\forall\, a, b, c \in L$: $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$.

Remark 5.13. Thus, if the user's prior knowledge is assumed to be in Π_m^+, a "no" answer to a monotone Boolean query always preserves the privacy of a "yes" answer to another monotone Boolean query. Roughly speaking, it is OK to disclose a negative fact while protecting a positive fact. This observation is especially helpful when A and B are given by query language expressions, whose monotonicity is often obvious.

5.3. PRIVACY FOR PRODUCT DISTRIBUTIONS. In this section, we shall study the problem of checking the privacy relation $\text{Safe}_\Pi(A, B)$ for sets $A, B \subseteq \Omega = \{0,1\}^n$ over the family $\Pi = \Pi_m^0$ of product distributions. The *independence* relation $A \perp_{\Pi_m^0} B$, defined by

$$A \perp_{\Pi_m^0} B \iff^{\text{def}} \forall P \in \Pi_m^0 : P[A]\,P[B] = P[AB],$$

has been studied by Miklau and Suciu, who proved the following necessary and sufficient criterion:

THEOREM 5.14 (MIKLAU & SUCIU). *For all $A, B \subseteq \Omega$, we have $A \perp_{\Pi_m^0} B$ if and only if sets A and B "share no critical coordinates," that is, when coordinates $1, 2, \ldots, n$ can be rearranged so that only $\omega[1], \omega[2], \ldots, \omega[k]$ determine if $\omega \in A$, and only $\omega[k+1], \omega[k+2], \ldots, \omega[k']$, where $k' \leqslant n$, determine if $\omega \in B$.*

PROOF. See Miklau and Suciu [2004]. □

Since $A \perp_{\Pi_m^0} B$ implies $\text{Safe}_{\Pi_m^0}(A, B)$, Miklau-Suciu criterion is a sufficient criterion for our notion of privacy. It is not a necessary one, even for $n = 2$: if we set $\Omega = \{00, 01, 10, 11\}$ and for $i = 1, 2$ define X_i by $(\omega \in X_i) \Leftrightarrow (\omega[i] = 1)$, then we have $\text{Safe}_{\Pi_m^0}(X_1, \bar{X}_1 \cup X_2)$ because for all $P \in \Pi_m^0$

$$P[X_1 \cap (\bar{X}_1 \cup X_2)] = P[X_1 \cap X_2] = P[X_1] \cdot P[X_2] \leqslant P[X_1] \cdot P[\bar{X}_1 \cup X_2],$$

but not $X_1 \perp_{\Pi_m^0} (\bar{X}_1 \cup X_2)$ since they share a critical coordinate #1.

Another sufficient criterion is given by Corollary 5.12, if we note that $\Pi_m^0 \subseteq \Pi_m^+$; it implies $\text{Safe}_{\Pi_m^0}(A, B)$ whenever A is an up-set and B is a down-set, or vice-versa. A little more generally, Proposition 5.11 implies

COROLLARY 5.15 (MONOTONICITY CRITERION). *Let $A, B \subseteq \Omega = \{0,1\}^n$. Relation $\text{Safe}_{\Pi_m^0}(A, B)$ holds if there exists a "mask" vector $z \in \Omega$ such that either one of the two conditions below is satisfied for $A_z = z \oplus A := \{z \oplus \omega \mid \omega \in A\}$ and $B_z = z \oplus B$:*

— $A_z B_z \wedge \bar{A}_z \bar{B}_z \subseteq A_z - B_z$ *and* $A_z B_z \vee \bar{A}_z \bar{B}_z \subseteq B_z - A_z$;
— $A_z B_z \vee \bar{A}_z \bar{B}_z \subseteq A_z - B_z$ *and* $A_z B_z \wedge \bar{A}_z \bar{B}_z \subseteq B_z - A_z$.

In particular, $\text{Safe}_{\Pi_m^0}(A, B)$ *holds if $z \oplus A$ is an up-set and $z \oplus B$ is a down-set.*

PROOF. By Proposition 5.11, either condition implies $\text{Safe}_{\Pi_m^+}(A_z, B_z)$, which in turn implies $\text{Safe}_{\Pi_m^0}(A_z, B_z)$. Finally, we have $\text{Safe}_{\Pi_m^0}(A_z, B_z) \Leftrightarrow \text{Safe}_{\Pi_m^0}(A, B)$ because the set of distributions $P(z \oplus \omega)$ over $\omega \in \Omega$ where $P \in \Pi_m^0$ is the same as Π_m^0 itself. □

It turns out that both the Miklau-Suciu and the monotonicity criteria are special cases of another simple yet surprisingly strong sufficient criterion for $\text{Safe}_{\Pi_m^0}(A, B)$. This sufficient criterion shall be called the *cancellation criterion*,

because its verification is equivalent to cancelling identical monomial terms in the algebraic expansion for the difference

$$P[A\bar{B}] \cdot P[\bar{A}B] \;-\; P[AB] \cdot P[\bar{A}\bar{B}], \tag{34}$$

where P is a product distribution written as in (25). Recall that expression (34) equals $P[A]P[B] - P[AB]$, see Proposition 3.8. In order to formulate the criterion in combinatorial (rather than algebraic) terms, we need the following definition:

Definition 5.16. The *pairwise matching function* $\text{match}(u, v)$ maps a pair (u, v) of vectors from $\Omega = \{0, 1\}^n$ to a single *match-vector* $w = \text{match}(u, v)$ in $\{0, 1, *\}^n$ as follows:

$$\forall\, i = 1 \ldots n: \quad w[i] = \begin{cases} u[i] & \text{if } u[i] = v[i]; \\ * & \text{if } u[i] \neq v[i]. \end{cases}$$

For example, pair $(01011, 01101)$ gets mapped into $01{*}{*}1$. We say that $v \in \Omega$ *refines* a match-vector w when v can be obtained from w by replacing its every star with a 0 or a 1. For every match-vector w, define the following two sets:

$$\begin{aligned} \text{Box}(w) &:= \{v \in \Omega \mid v \text{ refines } w\}; \\ \text{Circ}(w) &:= \{(u, v) \in \Omega \times \Omega \mid \text{match}(u, v) = w\}. \end{aligned}$$

Remark 5.17. Function "match" satisfies the following property: for all u, v, u', v' in $\{0, 1\}^n$, we have

$$\text{match}(u, v) = \text{match}(u', v') \;\Leftrightarrow\; u \wedge v = u' \wedge v' \;\&\; u \vee v = u' \vee v'. \tag{35}$$

Indeed, a coordinate that has the same bit-value in u and v stays the same in $u \wedge v$ and $u \vee v$, while a coordinate that is different in u versus v has value 0 in $u \wedge v$ and 1 in $u \vee v$. Hence, given $\text{match}(u, v)$, we can reconstruct both $u \wedge v$ and $u \vee v$ by replacing the $*$'s with 0's for $u \wedge v$ and with 1's for $u \vee v$; and vice-versa.

Now we are ready to state the cancellation criterion, which is a sufficient criterion for $\text{Safe}_{\Pi^0_m}(A, B)$, and also state a necessary criterion of a similar form, for comparison:

PROPOSITION 5.18 (CANCELLATION CRITERION). *For all $A, B \subseteq \Omega$, in order to establish* $\text{Safe}_{\Pi^0_m}(A, B)$ *it is sufficient to verify the following:*

$$\forall\, w \in \{0, 1, *\}^n: \quad |(AB \times \bar{A}\bar{B}) \cap \text{Circ}(w)| \;\leqslant\; |(A\bar{B} \times \bar{A}B) \cap \text{Circ}(w)|. \tag{36}$$

On the other hand, for all $A, B \subseteq \Omega$, if $\text{Safe}_{\Pi^0_m}(A, B)$ *holds, then:*

$$\forall\, w \in \{0, 1, *\}^n: \quad \left|(AB \times \bar{A}\bar{B}) \cap \text{Box}(w)^2\right| \;\leqslant\; |(A\bar{B} \times \bar{A}B) \cap \text{Box}(w)^2|. \tag{37}$$

Here $\text{Box}(w)^2$ *denotes* $\text{Box}(w) \times \text{Box}(w)$*, and* $|S|$ *denotes the size of set S.*

PROOF. Two subsets $S, S' \subseteq \text{Circ}(w)$ satisfy $|S| \leqslant |S'|$ if and only if there is an injective function that maps S into S'. As the preimages of $\text{match}(\cdot, \cdot)$ the sets $\text{Circ}(w)$ are all mutually disjoint and form a partition of $\Omega \times \Omega$. Hence, condition (36) is equivalent to the existence of an injective function F from $AB \times \bar{A}\bar{B}$ to $A\bar{B} \times \bar{A}B$ that maps each partition cell to itself, that is:

$$\forall\, u \in AB, \;\; \forall\, v \in \bar{A}\bar{B}: \quad \text{match}(u, v) \;=\; \text{match}(F(u, v)). \tag{38}$$

Suppose we have such an F, and let $P \in \Pi^0_{\mathrm{m}}$. By Proposition 5.2, since P is a product distribution, we have $P(\omega_1)P(\omega_2) = P(\omega_1 \wedge \omega_2)P(\omega_1 \vee \omega_2)$ for all $\omega_1, \omega_2 \in \Omega$, and therefore

$$\begin{aligned} P[A]\,P[B] \;-\; P[AB] \;&=\; P[A\bar{B}]\,P[\bar{A}B] \;-\; P[AB]\,P[\bar{A}\bar{B}] \quad \text{(Prop. 3.8)} \\ &= \sum_{\substack{\omega_1 \in A-B \\ \omega_2 \in B-A}} P(\omega_1)\,P(\omega_2) \;-\; \sum_{\substack{\omega'_1 \in AB \\ \omega'_2 \in \bar{A}\bar{B}}} P(\omega'_1)\,P(\omega'_2) \qquad (39) \\ &= \sum_{\substack{\omega_1 \in A-B \\ \omega_2 \in B-A}} P(\omega_1 \wedge \omega_2)\,P(\omega_1 \vee \omega_2) \;-\; \sum_{\substack{\omega'_1 \in AB \\ \omega'_2 \in \bar{A}\bar{B}}} P(\omega'_1 \wedge \omega'_2)\,P(\omega'_1 \vee \omega'_2). \end{aligned}$$

Every term in the right summation is canceled by an identical term in the left summation, with $(\omega_1, \omega_2) = F(\omega'_1, \omega'_2)$. The two terms are identical due to property (35). After the cancellation, we are left with a nonnegative expression, and that proves $\mathrm{Safe}_{\Pi^0_{\mathrm{m}}}(A, B)$.

To prove the necessary criterion (37), take some match-vector $w \in \{0, 1, *\}^n$ and consider the following product distribution defined as in Eq. (25) by its vector $(p_1, p_2, \ldots, p_n)$ of bit probabilities: $p_i = w[i]$ if $w[i] = 0$ or 1; $p_i = 1/2$ if $w[i] = *$. Then, for all vectors $v \in \mathrm{Box}(w)$ we have $P(v) = 1/2^k$ where $k =$ the number of stars in w; for all other vectors $P(v) = 0$. Therefore, $\forall\, S \subseteq \Omega$: $P[S] = 2^{-k} \cdot |S \cap \mathrm{Box}(w)|$, and inequality (37) is equivalent to $P[AB]\,P[\bar{A}\bar{B}] \leqslant P[A\bar{B}]\,P[\bar{A}B]$, which holds due to $\mathrm{Safe}_{\Pi^0_{\mathrm{m}}}(A, B)$. □

We hope that the combinatorial simplicity of the sufficient criterion given by Proposition 5.18 will allow highly scalable implementations that apply in real-life database auditing scenarios, where sets A and B are given via expressions in a query language. The theorems below justify our interest in the cancellation criterion:

THEOREM 5.19. *If sets A, B satisfy the Miklau-Suciu criterion, they also satisfy the cancellation criterion.*

PROOF. Assume that we have rearranged the coordinates so that only $\omega[1 \ldots k]$ determine if $\omega \in A$, and only $\omega[k+1 \ldots n]$ determine if $\omega \in B$ (see Theorem 5.14). To prove the cancellation condition (36), let us define an injective function F from $AB \times \bar{A}\bar{B}$ to $A\bar{B} \times \bar{A}B$ that satisfies the match-preservation property (38), as follows:

$$\begin{aligned} F(u, v) \;&=\; F(u[1 \ldots k]\,u[k+1 \ldots n],\; v[1 \ldots k]\,v[k+1 \ldots n]) \\ &:= \quad (u[1 \ldots k]\,v[k+1 \ldots n],\; v[1 \ldots k]\,u[k+1 \ldots n]). \end{aligned}$$

That is, function $F(u, v)$ swaps the last $n - k$ coordinates between the first and the second argument. The result of match(u, v) is the same as the result of match$(F(u, v))$ because, coordinate-wise, the same bits are matched. Therefore, F maps Circ(w) into itself, for every match-vector w.

Why does F map $AB \times \bar{A}\bar{B}$ into $A\bar{B} \times \bar{A}B$? Take any $u \in AB$ and $v \in \bar{A}\bar{B}$, and denote $(x, y) = F(u, v)$. The first k coordinates of x are the same as of u, therefore x belongs to A just like u does; the last $n - k$ coordinates of x are the same as of v, therefore x belongs to $\bar{B}$ just like v does. Analogously, the first k coordinates of y are the same as of v, so $y \notin A$, and the last $n - k$ coordinates of y are the same as of u, so $y \in B$. It follows that (x, y) in $A\bar{B} \times \bar{A}B$. □

THEOREM 5.20. *If sets A, B satisfy the monotonicity criterion, they also satisfy the cancellation criterion.*

Before we show that the cancellation criterion subsumes the monotonicity criterion, let us observe the following fact:

LEMMA 5.21. $\forall w \in \{0,1,*\}^n$, $\forall S \subseteq \mathrm{Circ}(w)$ *define*

$$\delta S \;:=\; \{(u \vee v', v \wedge u') \,|\, (u,v),\ (u',v') \in S\}. \tag{40}$$

Then we have: $\delta S \subseteq \mathrm{Circ}(w)$ *and* $|\delta S| \geqslant |S|$.

PROOF. First, let us prove $\delta S \subseteq \mathrm{Circ}(w)$ by showing that

$$(u,v),\ (u',v') \in \mathrm{Circ}(w) \;\Rightarrow\; (u \vee v',\ v \wedge u') \in \mathrm{Circ}(w).$$

Indeed, take some (u,v) and (u',v') in $\mathrm{Circ}(w)$, then by definition $\mathrm{match}(u,v) = \mathrm{match}(u',v') = w$. Let x be w with all stars replaced by 0, and y be w with all stars replaced by 1. As explained in Remark 5.17, we have $x = u \wedge v = u' \wedge v'$ and $y = u \vee v = u' \vee v'$. Then,

$$\begin{aligned}
(u \vee v') \wedge (v \wedge u') &= (u \wedge v \wedge u') \vee (v' \wedge v \wedge u') \\
&= (x \wedge u') \vee (x \wedge v) = x \vee x = x, \\
(u \vee v') \vee (v \wedge u') &= (u \vee v' \vee v) \wedge (u \vee v' \vee u') \\
&= (y \vee v') \wedge (y \vee u) = y \wedge y = y.
\end{aligned}$$

Again by the same reasoning as in Remark 5.17, the above equalities imply $\mathrm{match}(u \vee v',\ v \wedge u') = w$, and therefore $(u \vee v',\ v \wedge u') \in \mathrm{Circ}(w)$.

The proof of $|\delta S| \geqslant |S|$ is based on the Marica-Schönheim inequality [Marica and Schönheim 1969] (see also Section 19 in Bollobás [1986], and Aharoni and Holzman [1993]), which states that $\forall U \subseteq \{0,1\}^n$ and for operation $\omega - \omega' := \omega \wedge \neg\omega'$:

$$|\Delta U| \;\geqslant\; |U|, \quad \text{where } \Delta U \;:=\; \{\omega - \omega' \mid \omega,\ \omega' \in U\}.$$

Observe that in pairs $(u,v) \in \mathrm{Circ}(w)$ vector u can be computed from v by inverting the bits that correspond to stars in w. Therefore, we can replace all pairs in the subsets S and δS of $\mathrm{Circ}(w)$ by their second vectors, without change in the cardinality of these subsets. We can also discard all non-star (in w) coordinates, because they are the same in all vectors. Denote thus projected S and δS by $\widehat{S}$ and $\widehat{\delta S}$, and denote vectors u, v, u', v' without the non-star coordinates by $\hat{u}, \hat{v}, \hat{u}', \hat{v}'$. We have $\hat{u} = \neg\hat{v}$, $\hat{u}' = \neg\hat{v}'$, and:

$$\begin{aligned}
\widehat{\delta S} \;&=\; \{\hat{v} \wedge \hat{u}' \mid (u,v),\ (u',v') \in S\} \;=\; \{\hat{v} \wedge \neg\hat{v}' \mid \hat{v}, \hat{v}' \in \widehat{S}\} \\
&=\; \Delta\widehat{S}, \quad \text{implying } |\delta S| \;=\; |\widehat{\delta S}| \;=\; |\Delta\widehat{S}| \;\geqslant\; |\widehat{S}| \;=\; |S|. \quad \square
\end{aligned}$$

Having proven Lemma 5.21, we are now ready to prove Theorem 5.20:

PROOF (THEOREM 5.20). Let $A, B \subseteq \Omega = \{0,1\}^n$ be two sets that satisfy the monotonicity criterion (Corollary 5.15). Then, $\exists z \in \Omega$ such that sets $A_z = z \oplus A$ and $B_z = z \oplus B$ satisfy either one of the following two conditions:

— $A_z B_z \;\wedge\; \bar{A}_z \bar{B}_z \;\subseteq\; A_z - B_z$ and $A_z B_z \;\vee\; \bar{A}_z \bar{B}_z \;\subseteq\; B_z - A_z$;
— $A_z B_z \;\vee\; \bar{A}_z \bar{B}_z \;\subseteq\; A_z - B_z$ and $A_z B_z \;\wedge\; \bar{A}_z \bar{B}_z \;\subseteq\; B_z - A_z$.

We want to show that they satisfy the cancellation criterion (i.e., the sufficient criterion in Proposition 5.18).

First, note that, for every $z \in \Omega$, sets A and B satisfy the cancellation criterion if and only if sets $z \oplus A$ and $z \oplus B$ also do, because $(z, z) \oplus \mathrm{Circ}(w) = \mathrm{Circ}(z \oplus w)$. Therefore, we can assume that $z = 00\cdots0$ and ignore it. In what follows, we shall assume without loss of generality that $A_z = A$ and $B_z = B$ satisfy the second of the two conditions (if they satisfy the first, just swap the order of pairs in δS):

$$AB \vee \bar{A}\bar{B} \subseteq A\bar{B}, \quad AB \wedge \bar{A}\bar{B} \subseteq \bar{A}B. \tag{41}$$

Let us take an arbitrary match-vector $w \in \{0, 1, *\}^n$, define

$$S = (AB \times \bar{A}\bar{B}) \cap \mathrm{Circ}(w)$$

and show that $|S| \leqslant |(A\bar{B} \times \bar{A}B) \cap \mathrm{Circ}(w)|$. Indeed, by Lemma 5.21, for set δS defined in (40) we have $\delta S \subseteq \mathrm{Circ}(w)$ and $|S| \leqslant |\delta S|$. By (41), all pairs in δS are in $A\bar{B} \times \bar{A}B$: every pair has the form $(u \vee v', v \wedge u')$ where u and v belong to AB whereas u' and v' belong to $\bar{A}\bar{B}$. Therefore, (36) holds, and the cancellation criterion is satisfied. □

Remark 5.22. The sufficient condition in the cancellation criterion is not necessary. Here is a pair of sets that satisfies the privacy predicate $\mathrm{Safe}_{\Pi^0_m}(A, B)$, but does not satisfy the cancellation criterion:

$$A = \{011, 100, 110, 111\}; \quad B = \{010, 101, 110, 111\}.$$

Sets $(A-B) \times (B-A)$ and $AB \times \bar{A}\bar{B}$ can be conveniently represented in the form of a table:

$A-B$	$B-A$	match
100	010	$**0$
100	101	$10*$
011	010	$01*$
011	101	$**1$

match	AB	$\bar{A}\bar{B}$
$**0$	110	000
$***$	110	001
$***$	111	000
$**1$	111	001

We can see that $|A\bar{B} \times \bar{A}B \cap \mathrm{Circ}(***)| = 0$ and $|AB \times \bar{A}\bar{B} \cap \mathrm{Circ}(***)| = 2$ for these sets, violating (36). In the expression for $P[A]\,P[B] - P[AB]$, written as in (39), the product terms for the $**0$-matching pairs and for the $**1$-matching pairs cancel each other. The remaining terms result in expression

$$p_1^2 \cdot (1-p_2)^2 \cdot p_3(1-p_3) \;+\; (1-p_1)^2 \cdot p_2^2 \cdot p_3(1-p_3) \\ -\; 2 \cdot p_1(1-p_1) \cdot p_2(1-p_2) \cdot p_3(1-p_3),$$

which is nonnegative due to inequality $x^2 + y^2 \geqslant 2xy$.

6. *The Computational Complexity of Testing Safety*

We use techniques from multivariate polynomial optimization to test safety with respect to certain families Π of prior distributions on an agent's knowledge. Recall that a set $A \subseteq \Omega$ is Π-safe given $B \subseteq \Omega$ when for all distributions $P \in \Pi$, we have $P[A \mid B] \leqslant P[A]$, or equivalently, $P[AB] \leqslant P[A] \cdot P[B]$. As in some previous sections, we identify the set Ω of possible worlds with the hypercube $\{0, 1\}^n$.

For each $x \in \{0, 1\}^n$, we create variables $p_x \in [0, 1]$. We consider those families Π consisting of distributions $(p_x)_{x \in \{0,1\}^n}$ that can be described by a finite number

r of polynomial inequalities, together with the standard distribution equality and inequalities:

$$\alpha_1((p_x)_{x\in\{0,1\}^n}) \geqslant 0, \ldots, \alpha_r((p_x)_{x\in\{0,1\}^n}) \geqslant 0,$$
$$\textstyle\sum_{x\in\{0,1\}^n} p_x = 1, \quad \forall x \;\; p_x \geqslant 0.$$

We call such a family Π *algebraic*. For example, if we had the family of log-submodular distributions, then for all $x, y \in \{0,1\}^n$, we would have the constraint $p_x p_y - p_{x\wedge y} p_{x\vee y} \geqslant 0$. For the family of log-supermodular distributions, we would instead have $p_{x\wedge y} p_{x\vee y} - p_x p_y \geqslant 0$. Finally, for the family of product distributions, we would have both $p_x p_y - p_{x\wedge y} p_{x\vee y} \geqslant 0$ and $p_{x\wedge y} p_{x\vee y} - p_x p_y \geqslant 0$.

For sets A and B, and a family Π of distributions, we define the set $K(A, B, \Pi)$ of distributions $(p_x)_{x\in\{0,1\}^n}$ that satisfy

$$\sum_{w\in AB} p_w > \sum_{x\in A} p_x \sum_{y\in B} p_y$$
$$\alpha_1((p_x)_{x\in\{0,1\}^n}) \geqslant 0, \ldots, \alpha_r((p_x)_{x\in\{0,1\}^n}) \geqslant 0$$
$$\textstyle\sum_{x\in\{0,1\}^n} p_x = 1, \quad \forall x \;\; p_x \geqslant 0.$$

The following proposition is an equivalent algebraic formulation of the fact that in order for $\mathrm{Safe}_\Pi(A, B)$ to hold, there cannot be a single distribution $P \in \Pi$ for which $P[AB] > P[A] \cdot P[B]$. It follows immediately from the definition of $K(A, B, \Pi)$.

PROPOSITION 6.1. $\mathrm{Safe}_\Pi(A, B)$ *if and only if the set* $K(A, B, \Pi)$ *is empty.*

We are interested in algorithms that decide emptiness of $K(A, B, \Pi)$ in time polynomial or nearly polynomial in $N \stackrel{\text{def}}{=} 2^n$. Recall that n corresponds to the total number of possible records, and for a world $\omega \in \{0,1\}^n$, record i occurs in ω if and only if $\omega_i = 1$.

6.1. SPECIFIC DISTRIBUTIONS. In this section, we obtain efficient algorithms for testing safefty for certain interesting families Π of distributions.

We first obtain a necessary and sufficient condition for $A, B \subseteq \{0,1\}^n$ to be safe with respect to the family Π of product distributions by providing a deterministic algorithm. Its running time is $N^{O(\lg\lg N)}$, which is essentially polynomial for all practical purposes. The key observation is that while $K(A, B, \Pi)$ is $N = 2^n$-dimensional for general families of distributions, for product distributions it can be embedded into $\mathbb{R}^n$.

Indeed, it is easy to see that $K(A, B, \Pi)$ can be defined in variables $p_1, \ldots, p_n \in \mathbb{R}$ constrained by $p_i(1 - p_i) \geqslant 0$, and for which $P[AB] > P[A] \cdot P[B]$, where $P(\omega) = \prod_{i=1}^n p_i^{\omega[i]} \cdot (1 - p_i)^{1-\omega[i]}$ for all $\omega \in \{0,1\}^n$. We can write this with n variables and $n + 1$ inequalities. Notice that the inequality $P[AB] > P[A] \cdot P[B]$ can have an exponential number of terms in n. We apply the following simplified form of Theorem 3 of Basu et al. [1996]:

THEOREM 6.2. *Given a set* $K = \{\beta_1, \ldots, \beta_r\}$ *of* r *polynomials each of degree at most* d *in* s *variables with coefficients in* $\mathbb{R}$*, the problem of deciding whether there exist* $X_1, \ldots, X_s \in \mathbb{R}$ *for which* $\beta_1(X_1, \ldots, X_s) \geqslant 0, \ldots, \beta_r(X_1, \ldots, X_s) \geqslant 0$*, can be solved deterministically with* $\tau(rd)^{O(s)}$ *bit operations, where* τ *is the number of bits needed to describe a coefficient in* $\beta_1, \ldots, \beta_r$*.*

We apply this theorem to the set $K = K(A, B, \Pi)$. From the program above it is easy to see that τ, r, d, and s are all linear in n, and so emptiness (and hence safety) for product distributions can be decided in $n^{O(n)} = N^{O(\lg \lg N)}$ time.

The algorithm of Basu et al. uses sophisticated ideas from algebraic geometry over $\mathbb{R}$, and we cannot do it justice here. The general approach taken by such algorithms is to reduce a system of polynomial inequalities into a system of polynomial equalities by introducing slack variables, and then combining the multivariate polynomial equalities $p_i(x) = 0$ into a single equality $q(x) \stackrel{\text{def}}{=} \sum_x p_i^2(x) = 0$. One finds the critical points of $q(x)$, that is, the set $V_{\mathcal{C}}$ of common zeros of its partial derivatives over the complex field $\mathcal{C}$. By perturbing $q(x)$ and applying Bézout's Theorem, one can show that $|V_{\mathcal{C}}|$ is finite. Various approaches are used to find the subset $V_{\mathcal{R}}$ of $V_{\mathcal{C}}$ of real-valued points. Since $V_{\mathcal{R}}$ is finite, once it is found q is evaluated on each of its elements and the minimum value is taken. The main step is finding $V_{\mathcal{R}}$, and approaches based on Gröbner bases, resultant theory, and homotopy theory exist (see Parrilo and Sturmfels [2001]). The algorithm of Basu et al. [1996] may be practical. Indeed, a similar algorithm of Canny [1993] was implemented.

We remark that a simple trick allows one to further reduce the running time to $(|A| + |B|)^{O(\lg \lg(|A|+|B|))}$. First, observe that if either $|A|$ or $|B|$ is at least $\sqrt{N}$, then

$$N^{O(\lg \log N)} = (|A| + |B|)^{O(\lg \lg(|A|+|B|))},$$

and in this case we can simply run the algorithm above. Otherwise, we have that $|A| \cdot |B| < N$. Now, notice that the uniform distribution in which each $p_i = \frac{1}{2}$ is a product distribution. In order for $P[AB] \leqslant P[A]P[B]$ for this distribution, we need $\frac{|AB|}{N} \leqslant \frac{|A|\cdot|B|}{N^2}$, or equivalently, $N|AB| \leqslant |A| \cdot |B|$. If $AB \neq \emptyset$, then since $|A| \cdot |B| < N$ we cannot have $N|AB| \leqslant |A| \cdot |B|$. On the other hand, if $AB = \emptyset$, then $P[AB] \leqslant P[A]P[B]$ for any product distribution. It follows that if $|A| \cdot |B| < N$, testing safety reduces to testing whether or not A and B intersect, which can be done in poly$(|A| + |B|)$ time by a simple sorting algorithm. Thus, in all cases, the time complexity of testing safety for product distributions is $(|A| + |B|)^{O(\lg \lg(|A|+|B|))}$.

This approach generalizes to other algebraic families Π described by poly(n) constraints and $O(n)$ variables. For instance, a family of distributions for which $p_x = p_y$ whenever the Hamming weight of x and y are equal is described by $n + 1$ variables.

Even when the family Π of distributions requires N variables to describe, in certain cases we can obtain a polynomial-time algorithm for testing safety with respect to Π. Indeed, if the constraints α_i defining Π have degree at most 2 and there are only a constant number r of them, an algorithm in Grigoriev et al. [2003] shows how to decide emptiness of $K(A, B, \Pi)$ in $N^{O(r)}$ time. This algorithm makes black-box use of the earlier algorithm of Basu et al. [1996]. As an optimization, we note that if there are multiple linear equality constraints $L_i(X_1, \ldots, X_s) = 0$, it is helpful to combine them into a single quadratic constraint $\sum_i L_i^2 = 0$. This is because the running time is exponential in the number of constraints.

6.2. HARDNESS RESULTS. As the following theorem shows, even when the number N is not too large, we may need to restrict the class of distributions Π that we consider in order to efficiently test safety.

THEOREM 6.3. *If $NP \not\subseteq P/\text{poly}$,*[10] *then there is an algebraic Π for which the number of constraints is* poly(N)*, each constraint has degree at most* 2*, and for which deciding* $\text{Safe}_\Pi(A, B)$ *cannot be done in* poly(N) *time. This holds even if the deciding algorithm is allowed to preprocess the distribution Π with an unbounded amount of computational work, provided that the output of its preprocessing stage is a* poly(N)*-length bit string.*

Before proving the theorem, we first recall the optimization problem MAX-CUT. For an undirected unweighted graph $G = (V, E)$ on t vertices, a cut S of G is a set $S \subseteq V$ of vertices, and the cut size of the cut is the number of edges with one endpoint in S and the other in $V \setminus S$. A maximum cut is a cut of the largest possible cut size in G. The number of such edges is called the maximum cut size $\gamma(G)$, and the problem MAX-CUT is the problem of computing $\gamma(G)$. Note that one does not need to output a cut realizing $\gamma(G)$ to solve the MAX-CUT problem. However, given an oracle for computing $\gamma(\cdot)$, there is a standard reduction to obtain a maximum cut by iteratively deleting edges and checking whether they change the maximum cut size. Assuming $P \neq NP$, it is known [Karp 1972] that MAX-CUT cannot be solved in polynomial time.

We further restrict the MAX-CUT problem so that t is a power of 2. This is possible because we can increase the number of vertices of G by less than a factor of 2, so that now the number of vertices is a power of 2. If we make the vertices that we add be isolated vertices, then the maximum cut size of G remains the same.

Definition 6.4. The problem special MAX-CUT is the problem of determining whether $\gamma(G) > \frac{365t^2}{4608}$, given that the number t of vertices of G is a power of 2.

LEMMA 6.5. *Assuming $P \neq NP$,* special MAX-CUT *cannot be solved in* poly(t) *time.*

PROOF. Notice that MAX-CUT on graphs G' on $\frac{t}{4}$ vertices cannot be solved in poly(t)-time assuming $P \neq NP$. This is because if there were a poly(t)-time algorithm for solving MAX-CUT on graphs on $\frac{t}{4}$ vertices, the same algorithm would be a poly(t)-time algorithm for solving MAX-CUT on graphs on t vertices. It is not hard to show that any graph H on $\frac{t}{4}$ vertices satisfies $0 \leqslant \gamma(H) \leqslant \frac{t^2}{64}$, where the latter inequality is achieved by taking H to be a bipartite clique with $\frac{t}{8}$ vertices in each part. It is easy to see that $\frac{t^2}{64} < \frac{365t^2}{4608}$.

We need the fact that for every even integer $s \geqslant 2$ and non-negative integer $r \leqslant \frac{s^2}{4}$, there is a graph H_r on s vertices with $\gamma(H_r) = r$. This can be proven by induction on even integers s. It is clearly true for $s = 2$, since we can take H_0 to be the empty graph on 2 vertices, and H_1 to be a single edge. Suppose, inductively, that it is true for some value of $s \geqslant 2$. We want to show that for every $r \leqslant \frac{(s+2)^2}{4}$, there is a graph G_r on $s+2$ vertices with $\gamma(G_r) = r$. This clearly holds for $r \leqslant \frac{s^2}{4}$, since we can take G_r to be the disjoint union of H_r and 2 isolated vertices. For $r > \frac{s^2}{4}$, let S be a maximum cut of $H_{s^2/4}$. Denote the vertices of $H_{s^2/4}$ by V. Let u

[10] Recall that P/poly is the set of languages L for which there exists a polynomial-time algorithm A and an infinite advice sequence $(a_n)_{n\in\mathcal{N}}$ such that for every $x \in \{0,1\}^*$, $A(a_{|x|}, x) = 1$ if and only if $x \in L$.

and v be two vertices not in V. We connect u to a subset of vertices in S and v to a subset of vertices in $V \setminus S$ so that the total number of edges added is $r - \frac{s^2}{4}$. The cut $S \cup \{v\}$ is a maximum cut of the newly constructed graph G_r since all r edges in G_r participate in the cut. This works for all $r \leqslant \frac{s^2}{4} + s$. Notice that $\frac{(s+2)^2}{4} - \frac{s^2}{4} = s + 1$. If $r = \frac{s^2}{4} + s + 1$, then we also connect u to v, which increases the cut size of $S \cup \{v\}$ by one. This proves the inductive step.

Returning to the proof of the lemma, for each r with $0 \leqslant r \leqslant \frac{9t^2}{64}$, let J_r be a graph on $\frac{3t}{4}$ vertices with $\gamma(J_r) = r$.

Given a graph H on $\frac{t}{4}$ vertices, consider the graphs I_r on t vertices, where I_r is the disjoint union of H with J_r. Then $\gamma(I_r) = \gamma(H) + \gamma(J_r) = \gamma(H) + r$. Since $\gamma(H) \leqslant \frac{t^2}{64}$ and $\gamma(J_0) = 0$, we have that $\gamma(I_0) < \frac{365t^2}{4608}$. On the other hand, since $\gamma(J_{9t^2/64}) = \frac{9t^2}{64} > \frac{365t^2}{4608}$, we have that $\gamma(I_{9t^2/64}) > \frac{365t^2}{4608}$. Thus, there is some minimal value of r for which $\gamma(I_r) > \frac{365t^2}{4608}$. For this value of r, we have $\gamma(H) = \lceil \frac{365t^2}{4608} \rceil - r$. Thus, by solving special MAX-CUT for each graph I_r, we can determine $\gamma(H)$. It follows that special MAX-CUT cannot be solved in poly(t) time, if $P \neq NP$. □

We now prove Theorem 6.3.

PROOF OF THEOREM 6.3. Put $n = 1 + 2\log_2 t$, so that $N = 2t^2$. For each $u \in [t] \stackrel{\text{def}}{=} \{1, 2, \ldots, t\}$, associate an element $x_u \in \{0,1\}^n$. Associate each set $S = \{u, v\} \subseteq [t]$ of size 2 with a distinct element $y_S \in \{0,1\}^n$. Call such an S a 2-set. Let D_1 and D_2 be disjoint subsets of $\frac{t^2}{2}$ unassociated elements of $\{0,1\}^n$. Assume that 0^n is not in $D_1 \cup D_2$, and is not associated with any 2-set. Note that all of this is possible because $|D_1 \cup D_2| = t^2$ and the number of elements associated with a 2-set is $\binom{t}{2} \leqslant \frac{t^2}{2}$, while only t elements are associated with a value x_u. Thus, there are at least

$$N - \frac{3t^2}{2} - t = 2t^2 - \frac{3t^2}{2} - t > 0$$

unassociated elements of $\{0,1\}^n$ and not in $D_1 \cup D_2$ (for sufficiently large t).

We define Π by the following constraints. For each 2-set $S = \{u, v\}$, include the constraint:

$$p_{y_S} = 360 p_{x_u} \left(\frac{1}{20t} - p_{x_v} \right) + 360 p_{x_v} \left(\frac{1}{20t} - p_{x_u} \right).$$

For each $u \in [t]$, include the constraint:

$$0 = p_{x_u} \left(\frac{1}{20t} - p_{x_u} \right).$$

From this, we deduce that $p_{x_u} \in \{0, \frac{1}{20t}\}$. Moreover, we claim that $p_{y_S} \in \{0, \frac{9}{10t^2}\}$. To see this, note that for $S = \{u, v\}$, there are four cases: (1) $p_{x_u} = p_{x_v} = 0$, (2) $p_{x_u} = 0$ and $p_{x_v} = \frac{1}{20t}$, (3) $p_{x_u} = \frac{1}{20t}$ and $p_{x_v} = 0$, and (4) $p_{x_u} = p_{x_v} = \frac{1}{20t}$. We see that in cases (1) and (4), we have $p_{y_S} = 0$, while in cases (2) and (3) we have $p_{y_S} = \frac{9}{10t^2}$.

For each $z \in D_1 \cup D_2$, put $p_z = \frac{1}{2t^2}$. For those $z \notin D_1 \cup D_2 \cup \{0^n\}$ that are unassociated with any 2-set, put $p_z = 0$. Finally, put

$$p_{0^n} = 1 - \sum_{\text{2-sets } S=\{u,v\}} p_{y_S} - \sum_{u \in [t]} p_{x_u} - \sum_{z \in D_1 \cup D_2} p_z.$$

We thus have,

$$p_{0^n} \geqslant 1 - \binom{t}{2} \cdot \frac{9}{10t^2} - t \cdot \frac{1}{20t} - t^2 \cdot \frac{1}{2t^2} \geqslant 0.$$

Observe that $p_z \geqslant 0$ for all $z \in \{0, 1\}^n$ and $\sum_z p_z = 1$.

The constraints defining Π are equality constraints, which can each be converted into two inequality constraints. Observe that Π is algebraic and nonempty, the number of constraints is poly(N), and the constraints defining Π have degree at most 2. Moreover, each constraint can be described with $O(\log N)$ bits.

Given an input graph $G = ([t], E)$ and a parameter k, observe that the vertices $u \in [t]$ can be partitioned into two sets J and $[t] \setminus J$, where $u \in J$ if and only if $p_{x_u} = 0$. If $u \in [t] \setminus J$, then $p_{x_u} = \frac{1}{20t}$. Then, by the case analysis above, $p_{y_S} = \frac{9}{10t^2}$ if and only if one endpoint of e is in J and the other is in $[t] \setminus J$. Put $\bar{\gamma}(G) = \frac{\gamma(G)}{t^2}$. Put $m = \sum_{e \in E} p_{y_e}$. Hence, the maximum value of m is $\frac{9\gamma(G)}{10t^2} = \frac{9\bar{\gamma}(G)}{10}$.

We define query sets A and B as follows. Let $F_1 \subseteq D_1$ be an arbitrary subset of size $\frac{247t^2}{512}$, which is an integer for $t \geqslant 512$. Let $F_2 \subseteq D_2$ be an arbitrary subset of size $\frac{247t^2}{512}$.

Let $A = \cup_{e \in E} y_e \cup F_1$, and $B = \cup_{e \in E} y_e \cup F_2$. Then $AB = \cup_{e \in E} y_e$. Then, using that for each $z \in D_1 \cup D_2$ we have $p_z = \frac{1}{2t^2}$, the constraint $P[A]P[B] < P[AB]$ becomes

$$\left(m + \frac{247}{1024}\right)^2 < m, \tag{42}$$

since $P[A] = P[B] = P[\cup_{e \in E} y_e \cup F_2] = m + \frac{1}{2t^2} \cdot \frac{247t^2}{512} = m + \frac{247}{1024}$, and $P[AB] = m$.

The quadratic formula shows that this inequality holds if and only if

$$m \in \left(\frac{1}{2} - \frac{247}{1024} - \frac{3}{16}, \frac{1}{2} - \frac{247}{1024} + \frac{3}{16}\right).$$

We showed that $m \leqslant \frac{9\bar{\gamma}(G)}{10}$, and so if

$$\bar{\gamma}(G) \leqslant \frac{10}{9} \cdot \left(\frac{1}{2} - \frac{247}{1024} - \frac{3}{16}\right) = \frac{365}{4608},$$

then inequality (42) cannot hold.

We now turn to showing the converse, namely, that if $\bar{\gamma}G > \frac{365}{4608}$, then inequality (42) does hold for some distribution in Π. So suppose that $\bar{\gamma}(G) > \frac{365}{4608}$. If, also, $\bar{\gamma}(G) \leqslant \frac{10}{9} \cdot \left(\frac{1}{2} - \frac{247}{1024} + \frac{3}{16}\right)$, then by choosing the vertices in a maximum cut of G to be the set of vertices v for which $p_{x_v} = 0$, we have that $m \in \left(\frac{1}{2} - \frac{247}{1024} - \frac{3}{16}, \frac{1}{2} - \frac{247}{1024} + \frac{3}{16}\right)$, and so inequality (42) holds.

The only wrinkle comes when $\bar{\gamma}(G)$ is larger than $\frac{10}{9} \cdot (\frac{1}{2} - \frac{247}{1024} + \frac{3}{16})$. In this case, it suffices to exhibit a cut whose cut size lies in the interval

$$I = \left(\frac{10}{9}\left(\frac{t^2}{2} - \frac{247t^2}{1024} - \frac{3t^2}{16}\right),\ \frac{10}{9}\left(\frac{t^2}{2} - \frac{247t^2}{1024} + \frac{3t^2}{16}\right)\right).$$

By assumption on the maximum cut size, there is a cut S with cut size at least $\frac{10}{9} \cdot (\frac{t^2}{2} - \frac{247t^2}{1024} + \frac{3t^2}{16})$. Let $S = \{v_1, \ldots, v_r\}$. Consider the sequence of cuts $S_0 = S$, $S_1 = S \setminus \{v_1\}$, $S_2 = S \setminus \{v_1, v_2\}, \ldots, \emptyset$. The difference in cut sizes between consecutive cuts in this sequence is bounded by $t - 1$, the maximum degree of a vertex in G. Notice that the length of interval I is $\frac{10}{9} \cdot \frac{3t^2}{8} = \Omega(t^2)$. Since the last cut in the sequence, namely, $\emptyset$, has cut size 0, it follows that some cut in the sequence has cut size which is in interval I (for sufficiently large t). By the arguments above, it follows that $P[A]P[B] < P[AB]$.

It follows that $P[A]P[B] < P[AB]$ if and only if $\gamma(G) > \frac{10}{9} \cdot (\frac{t^2}{2} - \frac{247t^2}{1024} - \frac{3t^2}{16}) = \frac{365t^2}{4608}$. By Lemma 6.5, this cannot be solved in poly(t) = poly(N) time unless $P = NP$.

To prove the theorem, we must also allow the deciding algorithm access to a poly(N)-length bit string that does not depend on the query sets A and B. In this case, if Safe$_\Pi(A, B)$ could be decided in poly(N) time, then Special MAX-CUT on graphs containing t vertices could be solved in poly(t) time given a poly(t)-length bit string, and hence by the reduction in Lemma 6.5, MAX-CUT could also be solved in poly(t) time given a poly(t)-length bit string. But this implies there is a P/poly-algorithm for solving MAX-CUT, and since MAX-CUT is NP-complete, this would imply $NP \subseteq P$/poly. This contradicts the assumption of the theorem. □

6.3. HEURISTICS. For most families of distributions we will have to settle for a heuristic or an approximation for testing safety. If the program describing $K(A, B, \Pi)$ is multilinear (e.g., one can show this is the case for log-submodular and log-supermodular distributions), there are heuristics such as branch-and-bound or cutting-plane techniques. See page 2 of de Campos and Cozman [2005].

Here we describe the arguably most practical heuristic, the *sum-of-squares* heuristic, introduced in Shor [1987], Shor and Stetsyuk [1997], and Parrilo [2000], which works even for systems that are not multilinear. This heuristic was implemented with great success in Parrilo and Sturmfels [2001]. If $K(A, B, \Pi)$ is nonempty, that is Safe$_\Pi(A, B)$ does not hold, then the heuristic is guaranteed to report that $K(A, B, \Pi)$ is nonempty. On the other hand, there may be a false negative in the sense that if $K(A, B, \Pi)$ is empty, and so Safe$_\Pi(A, B)$ holds, then the heuristic may report that $K(A, B, \Pi)$ is nonempty, meaning that Safe$_\Pi(A, B)$ does not hold. One can reduce the likelihood of a false negative by increasing a parameter D given in the following description of the method.

The problem of minimizing a degree-d multivariate polynomial f over a set $K \subseteq \mathbb{R}^s$ is equivalent to finding the maximum $\gamma \in \mathbb{R}$ for which $f(x) - \gamma \geqslant 0$ for all $x \in K$. Let $\mathcal{P}_+^d(K)$ be the set of all polynomials in $\mathbb{R}[x_1, \ldots, x_s]$ of degree at most d which are nonnegative on every point in K. Thus, our problem is to find the maximum $\gamma \in \mathbb{R}$ for which $f - \gamma \in \mathcal{P}_+^d(K)$.

It is unknown how to optimize over $\mathcal{P}_+^d(K)$ efficiently, and so the following indirect route is taken. Define the set Σ^2:

$$\Sigma^2 = \left\{ f \in \mathbb{R}[x_1, \ldots, x_s] \;\middle|\; \exists g_1, \ldots, g_t \in \mathbb{R}[x_1, \ldots, x_s] \text{ s.t. } f = \sum_{i=1}^{t} g_i^2 \right\}.$$

Notice that Σ^2 is a subset of nonnegative polynomials, as every sum of squares of polynomials is nonnegative. It turns out that Σ^2 is in fact a strict subset of the non-negative polynomials, as shown non-constructively by Hilbert, and constructively by Motzkin who provided the polynomial

$$M(x, y, z) = x^4y^2 + x^2y^4 + z^6 - 3x^2y^2z^2.$$

Motzkin showed $M(x, y, z)$ is non-negative on $\mathbb{R}^3$, yet inexpressible as a sum of squares of polynomials. It turns out that every non-negative polynomial can be written as a sum of squares of rational functions (functions of the form $g_i(x)/h_i(x)$ for polynomials g_i and h_i), which was Hilbert's 17th problem, solved by Artin in 1927. While Σ^2 fails to capture all nonnegative polynomials, the following proposition is a compelling reason for studying it. The proposition is folklore, and is proven using semidefinite programming.

PROPOSITION 6.6. *For $f \in \mathbb{R}[x_1, \ldots, x_s]$ of bounded degree, the test "$f(x) \in \Sigma^2$" can be done in* poly(s) *time.*

Let $\Sigma^{2,d}$ be those $f(x) \in \Sigma^2$ of degree at most d. Then, $\Sigma^{2,d} \subseteq \mathcal{P}_+^d(\mathbb{R})$. To minimize $f(x)$ over $\mathbb{R}^s$, we find the largest $\lambda \in \mathbb{R}$ for which $f(x) - \lambda \in \Sigma^{2,d}$ via a binary search on λ and the proposition above. The value λ is a lower bound on $f(x)$ and in practice almost always agrees with the true minimum of f [Parrilo and Sturmfels 2001].

To minimize $f(x)$ over a set K constrained by polynomials, we need a few more tools. We could reduce the problem to minimizing a single polynomial, as mentioned in Section 6.1, but the following may work better in practice. We follow the presentation in Caramanis [2001].

Definition 6.7. The Algebraic Cone generated by elements $\beta_1, \ldots, \beta_t \in \mathbb{R}[x_1, \ldots, x_s]$ is the set

$$\mathcal{A}(\beta_1, \ldots, \beta_t) \stackrel{\text{def}}{=} \left\{ f \in \mathbb{R}[x_1, \ldots, x_l] \;\middle|\; f = \eta + \sum_{I \subseteq [t]} \eta_I \prod_{i \in I} \beta_i \right\},$$

where η and the η_I are in Σ^2, and $[t] = \{1, 2, \ldots, t\}$.

Thus, the algebraic cone can be thought of as the set of all affine combinations of all possible products of polynomials $\beta_1, \ldots, \beta_t$, where the coefficients of the affine combination are taken from Σ^2.

Definition 6.8. The Multiplicative Monoid $\mathcal{M}(\beta_1, \ldots, \beta_t)$ generated by $\beta_1, \ldots, \beta_t \in \mathbb{R}[x_1, \ldots, x_s]$ is the set of finite products of the β_i, including the empty product, which we set to 1.

The key result is a simplified form of the Positivstellensatz [Stengle 1974]:

THEOREM 6.9. *Given polynomials $\{f_1, \ldots, f_{t_1}\}, \{g_1, \ldots, g_{t_2}\}$ in $\mathbb{R}[x_1, \ldots, x_s]$, the set*

$$K \stackrel{\text{def}}{=} \{x \in \mathbb{R}^s : f_i(x) \geqslant 0, g_j(x) \neq 0, \forall i \in [t_1], j \in [t_2]\}$$

is empty if and only if $\exists F \in \mathcal{A}(f_1, \ldots, f_{t_1})$ *and* $G \in \mathcal{M}(g_1, \ldots, g_{t_2})$ *for which* $F + G^2$ *is the zero polynomial.*

Thus, for a set K described by f_i, and g_j of the form above, we consider $K' = K \cap \{x \in \mathbb{R}^s \mid \gamma - f(x) \geqslant 0, f(x) - \gamma \neq 0\}$. K' is empty if and only if $f(x) > \gamma$ for all $x \in K$.

Heuristics implemented in practice work by choosing a degree bound D, generating all $G \in \mathcal{M}(f - \gamma, g_1, \ldots, g_{t_2})$ of degree at most D (there are at most t_2^D such G), and checking if there is an $F \in \mathcal{A}(\gamma - f, f_1, \ldots, f_{t_1})$ for which $F + G^2 = 0$ via semidefinite programming. This is efficient for constant D, which usually suffices in practice. Better algorithms for special cases are based on alternative forms of the Positivstellensatz; see Putinar [1993] and Schmüdgen [1991].

7. *Conclusion*

We presented a novel approach to privacy where only gaining confidence in a sensitive fact is illegal, while losing confidence is allowed. We showed that this relaxation is significant and permits many more queries than with well-known approaches. In exchange, this gave us an opportunity to relax prior knowledge assumptions beyond current standards. Our hope is that work in this direction will help bridge the gap between theoretical soundness and practical usefulness of privacy frameworks.

One possible future goal is to obtain a better understanding of the families of sets and distributions that arise in practice, and to understand whether they admit efficient privacy tests. Another goal is to apply the new frameworks to online (proactive) auditing, which will require the modeling of a user's knowledge about the auditor's query-answering strategy.

ACKNOWLEDGMENTS. We thank Kenneth Clarkson for bringing our attention to the Four Functions Theorem, and Phokion Kolaitis for introducing us to co-clones mentioned in the proof of Proposition 5.6.

REFERENCES

ACKLEY, D. H., HINTON, G. E., AND SEJNOWSKI, T. J. 1985. A learning algorithm for Boltzmann machines. *Cognit. Sci. 9,* 1, 147–169.

AGRAWAL, R., BAYARDO, R. J., FALOUTSOS, C., KIERNAN, J., RANTZAU, R., AND SRIKANT, R. 2004. Auditing compliance with a hippocratic database. In *Proceedings of the 30th International Conference on Very Large Data Bases (VLDB'04)*. 516–527.

AGRAWAL, R., KIERNAN, J., SRIKANT, R., AND XU, Y. 2002. Hippocratic databases. In *Proceedings of the 28th International Conference on Very Large Data Bases (VLDB'02)*. 143–154.

AHARONI, R., AND HOLZMAN, R. 1993. Two and a half remarks on the Marica-Schönheim inequality. *J. London Math. Soc. 2-48,* 3, 385–395.

AHLSWEDE, R., AND DAYKIN, D. E. 1978. An inequality for the weights of two families of sets, their unions and intersections. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete 43*, 183–185.

AUSTRALIA. 1998. Australian privacy act of 1998. http://www.privacy.gov.au/ACT/privacyact.

BASU, S., POLLACK, R., AND ROY, M.-F. 1996. On the combinatorial and algebraic complexity of quantifier elimination. *J. ACM 43,* 6, 1002–1045.

BLUM, A., DWORK, C., MCSHERRY, F., AND NISSIM, K. 2005. Practical privacy: The SuLQ framework. In *Proceedings of the 24th ACM Symposium on Principles of Database Systems*. 128–138.

BÖHLER, E., CREIGNOU, N., REITH, S., AND VOLLMER, H. 2003. Playing with boolean blocks, part i: Posts lattice with applications to complexity theory. *ACM SIGACT News 34,* 4, 38–52 (Complexity Theory Column 42).

BOLLOBÁS, B. 1986. *Combinatorics: Set Systems, Hypergraphs, Families of Vectors and Combinatorial Probability*. Cambridge University Press.

CANADA. 2000. Personal information protection and electronic documents act. 2nd Session, 36th Parliament, 48-49 Elizabeth II, 1999–2000, Statutes of Canada.

CANNY, J. 1993. Improved algorithms for sign determination and existential quantifier elimination. *Comput. J. 36,* 5, 409–418. (Special Issue on Quantifier Elimination).

CARAMANIS, C. 2001. Non-convex optimization via real algebraic geometry. http://web.mit.edu/∼cmcaram/www/pubs/nonconvex_opt_review.pdf.

COVER, T. M., AND THOMAS, J. A. 2006. *Elements of Information Theory*. 2nd Ed. Wiley-Interscience, Chapter 12, 409–425.

CREIGNOU, N., KOLAITIS, P., AND ZANUTTINI, B. 2008. Structure identification of Boolean relations and plain bases for co-clones. *J. Comput. Syst. Sci. 74,* 7, 1103–1115.

DE CAMPOS, C. P., AND COZMAN, F. G. 2005. Computing lower and upper expectations under epistemic independence. In *Proceedings of the 4th International Symposium on Imprecise Probabilities and Their Applications*.

DINUR, I., AND NISSIM, K. 2003. Revealing information while preserving privacy. In *Proceedings of the 22nd ACM Symposium on Principles of Database Systems*. 202–210.

DWORK, C., AND NISSIM, K. 2004. Privacy-preserving datamining on vertically partitioned databases. In *Proceedings of the 24th International Conference on Cryptology (CRYPTO)*. 528–544.

E.U. PARLIAMENT. 1995. Directive 95/46/EC on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Official J. European Communities L. 281*, 31.

EVFIMIEVSKI, A., FAGIN, R., AND WOODRUFF, D. 2008. Epistemic privacy. In *Proceedings of the 27th ACM Symposium on Principles of Database Systems (PODS'08)*. 171–180.

EVFIMIEVSKI, A., GEHRKE, J., AND SRIKANT, R. 2003. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the 22nd ACM Symposium on Principles of Database Systems*. 211–222.

FAGIN, R., HALPERN, J. Y., MOSES, Y., AND VARDI, M. Y. 1995. *Reasoning About Knowledge*. The MIT Press, Cambridge, MA. (Paperbook edition appeared in 2001.)

FAGIN, R., HALPERN, J. Y., AND VARDI, M. Y. 1991. A model-theoretic analysis of knowledge. *J. ACM 91,* 2, 382–428.

FUJISHIGE, S. 2005. Submodular functions and optimization, *Annals of Discrete Mathematics*, vol. 58. Elsevier Science.

GRIGORIEV, D., DE KLERK, E., AND PASECHNIK, D. V. 2003. Finding optimum subject to few quadratic constraints in polynomial time. In *Proceedings of the Conference on Effective Methods in Algebraic Geometry (MEGA)*. Universität Kaiserslautern, Germany.

HINTIKKA, J. 1962. *Knowledge and Belief*. Cornell University Press, Ithaca, N.Y.

KARP, R. M. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*.

KENTHAPADI, K., MISHRA, N., AND NISSIM, K. 2005. Simulatable auditing. In *Proceedings of the 24th ACM Symposium on Principles of Database Systems*. 118–127.

KRIPKE, S. 1963. A semantical analysis of modal logic I: normal modal propositional calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 9*, 67–96.

LOVÁSZ, L. 1983. Submodular functions and convexity. In *Mathematical Programming—The State of the Art*, A. Bachem, M. Grötchel, and B. Korte, Eds. Springer-Verlag, 235–257.

MARICA, J., AND SCHÖNHEIM, J. 1969. Differences of sets and a problem of Graham. *Canadian Math. Bull. 12,* 5, 635–637.

MIKLAU, G., AND SUCIU, D. 2004. A formal analysis of information disclosure in data exchange. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. 575–586.

MOTWANI, R., NABAR, S. U., AND THOMAS, D. 2008. Auditing SQL queries. In *Proceedings of the IEEE 24th International Conference on Data Engineering*, 287–296. DOI: 10.1109/ICDE.2008.4497437.

NABAR, S. U., MARTHI, B., KENTHAPADI, K., MISHRA, N., AND MOTWANI, R. 2006. Towards robustness in query auditing. In *Proceedings of the 32nd International Conference on Very Large Data Bases*. 151–162.

PARRILO, P. A. 2000. Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization. Ph.D. dissertation, California Institute of Technology.

PARRILO, P. A., AND STURMFELS, B. 2001. Minimizing polynomial functions. In *Algorithmic and Quantitative Aspects of Real Algebraic Geometry in Mathematics and Computer Science*. 83–100.

PITAC. 2004. Revolutionizing health care through information technology. U.S. President's Information Technology Advisory Committee.

PÓLYA, G. 1954. *Mathematics and Plausible Reasoning, Volume I: Induction and Analogy in Mathematics* 1st Ed. Princeton University Press.
PÓLYA, G. 1957. *How to Solve It: A New Aspect of Mathematical Method* 2nd Ed. Princeton University Press. (Expanded ed. 2004.)
PÓLYA, G. 1968. *Mathematics and Plausible Reasoning, Volume II: Patterns of Plausible Inference* 2nd Ed. Princeton University Press.
PUTINAR, M. 1993. Positive polynomials on compact semi-algebraic sets. *Indiana Univ. Math J. 42*, 3.
SCHMÜDGEN, K. 1991. The k-moment problem for compact semialgebraic sets. *Ann. Math 289*, 203–206.
SHANNON, C. E. 1949. Communication theory of secrecy systems. *Bell System Tech. J. 28-4*, 656–715.
SHOR, N. Z. 1987. Class of global minimum bounds of polynomial functions. *Cybernetics 6*, 731–734.
SHOR, N. Z., AND STETSYUK, P. I. 1997. The use of a modification of the r-algorithm for finding the global minimum of polynomial functions. *Cybernetics Syst. Anal. 33*, 482–497.
STENGLE, G. 1974. A Nullstellensatz and a Positivstellensatz in semialgebraic geometry. *Annals of Math 207*, 87–97.
U. S. CONGRESS 1996. Health insurance portability and accountability act of 1996, United States public law 104–191. http://www.hhs.gov/ocr/hipaa.
VAN WRIGHT, G. H. 1951. *An Essay in Modal Logic*. North-Holland, Amsterdam.

RECEIVED DECEMBER 2008; REVISED APRIL 2010; ACCEPTED SEPTEMBER 2010