



HAL
open science

On the Additive Capacity Problem for Quantitative Information Flow

Konstantinos Chatzikokolakis

► **To cite this version:**

Konstantinos Chatzikokolakis. On the Additive Capacity Problem for Quantitative Information Flow. 15th International Conference on Quantitative Evaluation of SysTems (QEST 2018), Sep 2018, Beijing, China. pp.1-19. hal-01845330

HAL Id: hal-01845330

<https://inria.hal.science/hal-01845330v1>

Submitted on 20 Jul 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the Additive Capacity Problem for Quantitative Information Flow

Konstantinos Chatzikokolakis

CNRS, France

Abstract. Preventing information leakage is a fundamental goal in achieving confidentiality. In many practical scenarios, however, eliminating such leaks is impossible. It becomes then desirable to *quantify* the severity of such leaks and establish bounds on the threat they impose. Aiming at developing measures that are *robust* wrt a variety of operational conditions, a theory of channel *capacity* for the g -leakage model was developed in [1], providing solutions for several scenarios in both the multiplicative and the additive setting.

This paper continues this line of work by providing substantial improvements over the results of [1] for *additive* leakage. The main idea of employing the Kantorovich distance remains, but it is now applied to *quasimetrics*, and in particular the novel “*convex-separation*” quasimetric. The benefits are threefold: first, it allows to maximize leakage over a larger class of gain functions, most notably including the one of Shannon. Second, a solution is obtained to the problem of maximizing leakage over both priors and gain functions, left open in [1]. Third, it allows to establish an additive variant of the “Miracle” theorem from [3].

Keywords: Quantitative information flow · capacity · Kantorovich distance.

1 Introduction

Preventing sensitive information from being leaked is a fundamental goal of computer security. There are many situations, however, in which completely eliminating such leaks is impossible for a variety of reasons. Sometimes the leak is *intentional*: we *want* to extract knowledge from a statistical database; sometimes it is due to *side channels* that are hard or impossible to fully control; sometimes the leak is in exchange to a *service*, as in the case of Location-Based Services; sometimes it is in exchange for *efficiency*: i.e. using a weaker but more efficient anonymous communication system.

In these cases, it becomes crucial to *quantify* such leaks, measure how important the threat they pose is and decide whether they can be tolerated or not. This problem is studied in the area of *quantitative information flow*, in which much progress has been done in recent years, both from a foundational viewpoint [17, 21, 11, 24, 22, 12, 6, 3], but also in the development of counter-measures and verification techniques [4, 5, 10, 19, 27, 23, 16, 8, 7, 25], and the analysis of real systems [14, 20, 15, 18].

Robustness is a fundamental theme in this area; we aim at developing measures and bounds that are robust wrt a variety of adversaries and operational scenarios. In the context of the successful g -leakage model, the operational scenario is captured by a gain function g , and the adversary’s knowledge by a prior π . Developing the theme of

robustness in this model, [1] studied the theory of *channel capacity*, that is the problem of maximizing leakage over π for a fixed g , maximizing over g for a fixed π , or maximizing over both π and g . Comparing the system’s prior and posterior vulnerability can be done either *multiplicatively* or *additively*, leading to a total of six capacity scenarios.

In this paper we make substantial progress in two of the scenarios for additive leakage, namely in maximizing over g alone, or over both π, g . When maximizing over g , we quickly realize that if we allow vulnerability to take values in the whole $\mathbb{R}_{\geq 0}$, we can always *scale* it up, leading to unbounded capacity. In practice, however, it is common to measure vulnerability within a predefined range; for instance, vulnerabilities capturing the probability of some unfortunate event (e.g. Bayes vulnerability) take values in $[0, 1]$, while vulnerabilities measuring bits of information (e.g. Shannon vulnerability) take values in $[0, \log_2 |\mathcal{X}|]$. It is thus natural to restrict to a *class* \mathcal{G} of gain functions, in which the range of vulnerabilities is limited. In [1], this is achieved by the class $\mathbb{G}^1\mathcal{X}$ of *1-spanning* gain functions, in which the gain of different secrets varies by at most 1.

Although $\mathbb{G}^1\mathcal{X}$ provides a solution for capacity, this choice is not completely satisfactory from the point of view of robustness, since it excludes important vulnerability functions. Most notably, *Shannon vulnerability* (the complement of entropy) is not k -spanning for *any* k , hence the capacity bound for $\mathbb{G}^1\mathcal{X}$ does not apply, and indeed the leakage in this case (known as *mutual information*) does *exceed* the bound. In this paper we take a more permissive approach, by imposing the 1-spanning condition not on g itself, but on the corresponding vulnerability function V_g , leading to the class $\mathbb{G}^\dagger\mathcal{X}$. Since *any* vulnerability is k -spanning for some k , this class does not a priori exclude any type of adversary, it only restricts the range of values.

Solving the capacity problem for $\mathbb{G}^\dagger\mathcal{X}$ is however not straightforward. It turns out that the core technique from [1], namely the use of the *Kantorovich distance* on the *hyper-distribution* produced by the channel, can still be applied. However, substantial modifications are needed, involving the use of *quasimetrics*, and in particular the novel “*convex-separation*” quasimetric, replacing the total variation used in [1]. These improvements not only lead to a solution to the problem of maximizing leakage over $g : \mathbb{G}^\dagger\mathcal{X}$, but also lead to a solution for the third scenario of maximizing over both π, g , as well as to a variant of the “Miracle” theorem for the additive setting.

In detail, the paper makes the following contributions to the study of g -capacity:

- We present a general technique for computing additive capacity wrt a class of gain functions \mathcal{G} , using the Kantorovich distance over a properly constructed quasimetric.
- This technique is instantiated for $\mathbb{G}^1\mathcal{X}$ using the total variation metric, recovering the results of [1] in a more structured way.
- The same technique is then instantiated for the larger class $\mathbb{G}^\dagger\mathcal{X}$, using the novel “convex-separation” quasimetric for which an efficient solution is provided.
- The results for $\mathbb{G}^\dagger\mathcal{X}$ also provide an immediate solution to the scenario of maximizing over both π, g , which was left completely open in [1].
- Finally, the results for $\mathbb{G}^\dagger\mathcal{X}$ lead to an “Additive Miracle” theorem, similar in nature to the “Miracle” theorem of [3] for the multiplicative case.

Acknowledgements All results were obtained in the process of preparing a manuscript on Quantitative Information Flow with my long-term collaborators M. Alvim, C. Morgan, A. McIver, C. Palamidessi and G. Smith, and were heavily influenced by their feedback.

$$\begin{array}{c|cccc} \pi & C & y_1 & y_2 & y_3 & y_4 \\ \hline 1/3 & x_1 & 1 & 0 & 0 & 0 \\ 1/3 & x_2 & 0 & 1/2 & 1/4 & 1/4 \\ 1/3 & x_3 & 1/2 & 1/3 & 1/6 & 0 \end{array} \longrightarrow \begin{array}{c|cccc} J & y_1 & y_2 & y_3 & y_4 \\ \hline x_1 & 1/3 & 0 & 0 & 0 \\ x_2 & 0 & 1/6 & 1/12 & 1/12 \\ x_3 & 1/6 & 1/9 & 1/18 & 0 \end{array} \longrightarrow \begin{array}{c|ccc} [\pi \triangleright C] & 1/2 & 5/12 & 1/12 \\ \hline x_1 & 2/3 & 0 & 0 \\ x_2 & 0 & 3/5 & 1 \\ x_3 & 1/3 & 2/5 & 0 \end{array}$$

Fig. 1. A prior π (type $\mathbb{D}\mathcal{X}$), a channel C (rows have type $\mathbb{D}\mathcal{Y}$), a joint J (type $\mathbb{D}(\mathcal{X} \times \mathcal{Y})$), and a hyper $[\pi \triangleright C]$ (type $\mathbb{D}^2\mathcal{X}$, each column has type $\mathbb{D}\mathcal{X}$).

2 Preliminaries

Channels and their effect on the adversary’s knowledge A *channel* C is a simple probabilistic model describing the behavior of a system that takes input values from a finite set \mathcal{X} (the secrets) and produces outputs from a finite set \mathcal{Y} (the observations). Formally, it is a *stochastic* $|\mathcal{X}| \times |\mathcal{Y}|$ *matrix*, meaning that elements are non-negative and rows sum to 1. $C_{x,y}$ can be thought of as the conditional probability of producing y when the input is x .

We denote by $\mathbb{D}\mathcal{A}$ the set of all discrete distributions on \mathcal{A} , and by $[a]:\mathbb{D}\mathcal{A}$ the *point* distribution, assigning probability 1 to $a:\mathcal{A}$. Given C and a distribution $\pi:\mathbb{D}\mathcal{X}$, called the *prior*, we can create a *joint* distribution $J:\mathbb{D}(\mathcal{X} \times \mathcal{Y})$ as $J_{x,y} = \pi_x C_{x,y}$. When J is understood, it is often written in the usual notation $p(x, y)$, in which case the conditional probabilities $p(y|x) = p(x,y)/p(x)$ coincide with $C_{x,y}$ (when $p(x)$ is non-zero) and the x -marginals $p(x) = \sum_y p(x, y)$ coincide with π_x .

The prior π can be thought of as the *initial knowledge* that the adversary has about the secret. When secrets are passwords, for instance, she might know that some are more likely to be chosen than others. Always assuming that C is known to the adversary, each output y provides evidence that allows her to update her knowledge, creating a *posterior* distribution δ^y , defined as $\delta_x^y = p(x,y)/p(y)$. This, of course, can be done for each output; every $y:\mathcal{Y}$ potentially provides information to the adversary leading to an updated probabilistic knowledge δ^y . But not all outputs have the same status; each happens with a different marginal probability $p(y) = \sum_x p(x, y)$, denoted by a_y .

Hence, the *effect of a channel* C to the adversary’s prior knowledge π , is to produce a set of posteriors δ^y , each with probability a_y . It is conceptually useful to view this outcome as a single *distribution on distributions*, called a *hyper*-distribution or just *hyper*. Such a hyper has type $\mathbb{D}^2\mathcal{X}$ and is denoted by $[\pi \triangleright C] = \sum_y a_y [\delta^y]$.¹ The a_y -component of $[\pi \triangleright C]$ is called the *outer* distribution, expressing the probability of obtaining the posteriors δ^y , called the *inner* distributions.

An example of all constructions is given in Fig. 1. From a channel C and the uniform prior π we can construct the joint J by multiplying (element-wise) each column of C by π . J is then a single distribution assigning probabilities to each pair (x, y) . To construct the hyper $[\pi \triangleright C]$, we normalize (i.e. divide by $p(y)$) each column $J_{-,y}$, forming the posterior δ^y . The marginals $p(y)$ become the outer probabilities a_y , labeling the columns of $[\pi \triangleright C]$. Finally, note that $[\pi \triangleright C]$ no longer records the original label y of each column. As a consequence, the columns y_2 and y_3 , both producing the *same posterior*

¹ The notation is due to the fact that distributions can be convexly combined ($\mathbb{D}\mathcal{A}$ is a vector space). $\sum_y a_y [\delta^y]$ is exactly the hyper assigning probability a_y to each δ^y .

$\delta^{y_2} = \delta^{y_3} = (0, 3/5, 2/5)$, are merged in $[\pi \triangleright C]$, which assigns to that posterior the combined probability $p(y_2) + p(y_3) = 5/12$. This phenomenon happens automatically by the construction of the hyper.

Vulnerability and leakage A fundamental notion in measuring the information leakage of a system is that of a *vulnerability function* $V : \mathbb{D}\mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$. The goal of $V(\pi)$ is to measure how vulnerable a system is when the adversary has knowledge π about the secret. To create a suitable vulnerability function we need to consider the operational scenario at hand: we first determine what the adversary is trying to achieve, then take $V(\pi)$ as a measure of how successful the adversary is in that goal. Clearly, no single function can capture all operational scenarios; as a consequence a variety of vulnerability functions has been proposed in the literature, each having a different operational interpretation.

For instance, *Bayes-vulnerability* $V_B(\pi) := \max_x \pi_x$ measures the probability of success of an adversary who tries to guess the complete secret in one try; Shannon-vulnerability $V_H(\pi) := \log_2 |\mathcal{X}| + \sum_x \pi_x \log_2 \pi_x$ (the complement of entropy) measures the expected number of Boolean questions needed to reveal the secret; and Guessing-vulnerability $V_G(\pi) = |\mathcal{X}|^{1/2} - \sum_i i \pi_i$ (where the i -indexing of \mathcal{X} is in non-decreasing probability order) measures the expected numbers of tries to guess the secret correctly.

To study vulnerability in a unifying way, the g -leakage framework was introduced in [3], in which the operating scenario is parametrized by a (possibly infinite) set of *actions* \mathcal{W} , and a *gain function* $g : \mathcal{W} \times \mathcal{X} \rightarrow \mathbb{R}$. Intuitively, \mathcal{W} consists of actions that the adversary can perform to *exploit* his knowledge about the system. Then, $g(w, x)$ models the adversary's reward when he performs the action w and the actual secret is x . In such an operational scenario, it is natural to define g -vulnerability V_g as the expected gain of a rational adversary who chooses the best available action:

$$V_g(\pi) := \sup_w \sum_x \pi_x g(w, x).$$

The g -leakage framework is quite expressive, allowing to obtain a variety of vulnerability functions as special cases for suitable choices of g . For instance, by picking $\mathcal{W} = \mathcal{X}$ and the *identity* gain function given by $g_{\text{id}}(w, x) = 1$ iff $w = x$ and 0 otherwise, we get $V_{g_{\text{id}}} = V_B$. Similarly, we can construct gain functions expressing Shannon (for which an infinite \mathcal{W} is needed) and Guessing vulnerabilities, as well as a variety of other operational scenarios. In fact, it can be shown that any continuous and convex vulnerability $V : \mathbb{D}\mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ can be written as V_g for some g [2].

For expressiveness, it is crucial to allow g to potentially take negative values, and \mathcal{W} to be infinite. However, it is desirable that V_g *itself* be non-negative and finite-valued, since it is meant to express vulnerability. As a consequence we always restrict to the class of $\mathbb{G}\mathcal{X}$ of gain functions, defined as those such that $V_g : \mathbb{D}\mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$. As already discussed in the introduction, it is often desirable to further restrict to subsets of $\mathbb{G}\mathcal{X}$.

Having established a way to measure vulnerability in the prior case, we move on to measuring how vulnerable our system is after observing the output of a channel C . Viewing the outcome of C on π as the hyper $[\pi \triangleright C]$, there is a natural answer: we can measure the vulnerability of each posterior (inner) distribution of $[\pi \triangleright C]$, then average by the outer probabilities, leading to the following definition of *posterior g -vulnerability*:

$$V_g[\pi \triangleright C] := \sum_y a_y V_g(\delta^y) \quad \text{where} \quad [\pi \triangleright C] = \sum_y a_y [\delta^y]$$

Finally, information *leakage* is measured by comparing the vulnerability in the prior and posterior case. Depending on how we compare the two vulnerabilities, this leads to the *additive* or *multiplicative* leakage, defined as:

$$\mathcal{L}_g^+(\pi, C) := V_g[\pi \triangleright C] - V_g(\pi), \quad \mathcal{L}_g^\times(\pi, C) := \frac{V_g[\pi \triangleright C]}{V_g(\pi)}.$$

Additive g -capacities A fundamental theme when measuring information leakage is *robustness*; we need bounds that are robust wrt a variety of different adversaries and operational scenarios. Following this theme, since the g -leakage of a channel C depends on both the prior π and the gain function g , it is natural to ask what is the *maximum leakage* of C , over a *class* of gain functions $\mathcal{G} \subseteq \mathbb{G}\mathcal{X}$ and a class of priors $\mathcal{D} \subseteq \mathbb{D}\mathcal{X}$. This maximum leakage is known as the *capacity* of C .

Definition 1. *The additive $(\mathcal{G}, \mathcal{D})$ -capacity of C , for $\mathcal{G} \subseteq \mathbb{G}\mathcal{X}$, $\mathcal{D} \subseteq \mathbb{D}\mathcal{X}$, is*

$$\mathcal{ML}_{\mathcal{G}}^+(\mathcal{D}, C) := \sup_{g: \mathcal{G}, \pi: \mathcal{D}} \mathcal{L}_g^+(\pi, C).$$

For brevity, when maximizing *only* over π for a *fixed* g , we write $\mathcal{ML}_g^+(\mathcal{D}, C)$ instead of $\mathcal{ML}_{\{g\}}^+(\mathcal{D}, C)$; similarly, when π is fixed we write $\mathcal{ML}_{\mathcal{G}}^+(\pi, C)$; for specific classes, say $\mathcal{G} = \mathbb{G}\mathcal{X}$, $\mathcal{D} = \mathbb{D}\mathcal{X}$, we write $\mathcal{ML}_{\mathbb{G}}^+(\mathbb{D}, C)$ instead of $\mathcal{ML}_{\mathbb{G}\mathcal{X}}^+(\mathbb{D}\mathcal{X}, C)$. We can maximize over π , or g , or both, getting three scenarios for additive capacity. The multiplicative capacity $\mathcal{ML}_{\mathcal{G}}^\times(\mathcal{D}, C)$, defined similarly, is outside the scope of this paper; the corresponding three scenarios are studied in [1].

For the first scenario, g is fixed and we maximize over the whole $\mathbb{D}\mathcal{X}$. For some gain functions an efficient solution exists; for instance, for g_H giving Shannon vulnerability, $\mathcal{ML}_{g_H}^+(\mathbb{D}, C)$ is the Shannon capacity (maximum transmission rate)² which can be computed using the well-known Blahut-Arimoto algorithm [13]. However, for g_{id} (giving Bayes vulnerability), bounding $\mathcal{ML}_{g_{id}}^+(\mathbb{D}, C)$ is known to be NP-complete [1], which of course leaves no hope for a general solution.

The second scenario (fixed π , maximize over g) is the main focus of this paper and is studied in detail in §3. Our solution turns out to also provide an answer for the third scenario (maximize over both π, g), discussed in §3.5.

3 Computing additive capacities

In this section we study the problem of computing the additive (\mathcal{G}, π) -capacity. We quickly realize, however, that the unrestricted maximization over the whole $\mathbb{G}\mathcal{X}$ yields unbounded leakage. The problem is the *unbounded range* of V_g , and can be illustrated by “scaling” g . Define the scaling of g by $k > 0$ as $g_{\times k}(w, x) = k \cdot g(w, x)$. It is easy to show [1] that this operation gives leakage $\mathcal{L}_{g_{\times k}}^+(\pi, C) = k \cdot \mathcal{L}_g^+(\pi, C)$, and since k can be arbitrary we get that $\mathcal{ML}_{\mathbb{G}}^+(\pi, C) = +\infty$.³

² Which is why we generally refer to the maximization of leakage as “capacity”.

³ The same phenomenon happens for *multiplicative* leakage, this time demonstrated by shifting. To keep $\mathcal{ML}_{\mathcal{G}}^\times(\pi, C)$ bounded we can restrict to the class $\mathbb{G}^+\mathcal{X}$ of *non-negative* gain functions.

There are important classes of gain functions, however, which effectively *limit* the range of V_g , causing the additive leakage to remain bounded. Even when $\mathcal{ML}_{\mathcal{G}}^+(\pi, C)$ is finite, computing it efficiently is non-trivial. A solution can be obtained by exploiting the fact that $\mathcal{ML}_{\mathcal{G}}^+(\pi, C)$ is connected to the well-known *Kantorovich distance* between $[\pi]$ and $[\pi \triangleright C]$.

This section proceeds as follows. In §3.1 we recall the Kantorovich distance and then use it in §3.2 to obtain a generic technique for computing $\mathcal{ML}_{\mathcal{G}}^+(\pi, C)$, in time linear on the size of C , using properties of \mathcal{G} . We apply these bounds to obtain efficient solutions for the class $\mathbb{G}^1\mathcal{X}$ of 1-spanning gain functions in §3.3, as well as the class $\mathbb{G}^\dagger\mathcal{X}$ of gain functions giving 1-spanning vulnerability in §3.4. Finally, §3.5 discusses the scenario of maximizing over both π and g .

3.1 The Kantorovich and Wasserstein distances

We begin by recalling the Kantorovich distance between probability distributions. Given $\alpha: \mathbb{D}\mathcal{A}$ and random variable $F: \mathcal{A} \rightarrow \mathbb{R}$ we write $\mathcal{E}_\alpha F$ for the expected value of F over α , or $\mathcal{E}_{x \sim \alpha} F(x)$ to make precise the variable we are averaging over. Observe that for a point distribution centered at $a \in \mathcal{A}$ we have $\mathcal{E}_{[a]} F = F(a)$.

A function $d: \mathcal{A}^2 \rightarrow \mathbb{R}$ is called a *quasimetric* iff it satisfies the triangle inequality and $d(a, a') = 0 \wedge d(a', a) = 0$ iff $a = a'$. If d is also symmetric it is called a *metric*. The set of all quasimetrics on \mathcal{A} is denoted by $\mathbb{M}\mathcal{A}$. Although less frequently used than metrics, quasimetrics will play an important role in computing additive capacity in §3.4.

A natural quasimetric on \mathbb{R} is given by

$$d_{\mathbb{R}}^{\leq}(x, y) := \max\{y - x, 0\}.$$

Intuitively, $d_{\mathbb{R}}^{\leq}(x, y)$ measures “how much smaller” than y is x ; 0 means that x is no smaller than y . This quasimetric can be extended to $x, y \in \mathbb{R}^n$ as $d_{\mathbb{R}^n}^{\leq}(x, y) = \sum_i d_{\mathbb{R}}^{\leq}(x_i, y_i)$, giving an “asymmetric Manhattan” distance.

A function $F: \mathcal{A} \rightarrow \mathbb{R}$ is called d, d_T -Lipschitz iff

$$d_T(F(a), F(a')) \leq d(a, a') \quad \text{for all } a, a' \in \mathcal{A}. \quad (1)$$

The set of all such functions (also called *contractions*) is denoted by $\mathbb{C}^{d, d_T} \mathcal{A}$. For the source metric, a scaled distance $k \cdot d$ (for some $k \geq 0$) is often used. For the target metric d_T the Euclidean distance $d_{\mathbb{R}}$ is commonly employed (in which case we might simply write d -Lipschitz for brevity). In this section, however, we consider functions that are $d, d_{\mathbb{R}}^{\leq}$ -Lipschitz, which holds iff

$$F(a') - F(a) \leq d(a, a') \quad \text{for all } a, a' \in \mathcal{A}. \quad (2)$$

Note that the max from the definition of $d_{\mathbb{R}}^{\leq}$ is not needed, since $d(a, a')$ is non-negative.

The Kantorovich construction allows us to *lift* a metric d on \mathcal{A} to a metric on probability distributions over \mathcal{A} . The standard construction is done by maximizing $|\mathcal{E}_\alpha F - \mathcal{E}_{\alpha'} F'|$ over all functions that are $d, d_{\mathbb{R}}^{\leq}$ -Lipschitz. Note that the Euclidean distance is implicitly used twice in this construction: first, in the Lipschitz condition and second, for comparing $\mathcal{E}_\alpha F$ and $\mathcal{E}_{\alpha'} F'$. We can, however, define variants of the Kantorovich by using any other distance on \mathbb{R} . Our purpose is to work with quasimetrics, hence we employ $d_{\mathbb{R}}^{\leq}$, leading to the following definition.

Definition 2. *The Kantorovich quasimetric is the mapping $\mathbb{K}^< : \mathbb{M}\mathcal{A} \rightarrow \mathbb{M}\mathbb{D}\mathcal{A}$:*

$$\mathbb{K}^<(d)(\alpha, \alpha') := \sup_{F: \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathcal{A}} d_{\mathbb{R}}^{\leq}(\mathcal{E}_{\alpha} F, \mathcal{E}_{\alpha'} F) = \sup_{F: \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathcal{A}} \mathcal{E}_{\alpha'} F - \mathcal{E}_{\alpha} F .$$

Note that that max was again dropped from $d_{\mathbb{R}}^{\leq}$, since the sup is anyway non-negative ($\mathcal{E}_{\alpha'} F - \mathcal{E}_{\alpha} F = 0$ for F constant).

An important property of the Kantorovich distance is that it has a dual formulation as the Wasserstein (or “earth-moving”) metric, for which efficient algorithms exist. Earth moving measures measuring the cost of transforming one distribution into another, using the underlying distance d as the cost function. Given two distributions $\alpha, \alpha' : \mathbb{D}\mathcal{A}$ (the “source” and “target”), an earth-moving *strategy* is a joint distribution $S \in \mathbb{D}\mathcal{A}^2$ whose two marginals are α and α' . We write $\mathcal{S}_{\alpha, \alpha'}$ for the set of such strategies. The Wasserstein distance is then defined as the minimum transportation cost; similarly to Kantorovich, it provides a lifting from $\mathbb{M}\mathcal{A}$ to $\mathbb{M}\mathbb{D}\mathcal{A}$.

Definition 3. *The Wasserstein distance is the mapping $\mathbb{W} : \mathbb{M}\mathcal{A} \rightarrow \mathbb{M}\mathbb{D}\mathcal{A}$ given by:*

$$\mathbb{W}(d)(\alpha, \alpha') := \inf_{S: \mathcal{S}_{\alpha, \alpha'}} \mathcal{E}_S d .$$

The well-known Kantorovich-Rubinstein theorem [26] states that, if (d, \mathcal{A}) is a *separable* metric space then $\mathbb{K}(d) = \mathbb{W}(d)$. For our purposes, we will use this result in the restricted case where one of the two distributions is a *point* distribution $[a]$. This restriction is useful for two reasons: first, it allows us to give a simplified proof, adjusted to our $\mathbb{K}^<$ variant, drop the assumption of separability and allow d to be a quasimetric. Second, we show that that $\mathbb{W}(d)([a], \alpha)$ can be easily obtained as the expected (wrt α) distance between a and the elements in the support of α .

We first fix some notation: given $d \in \mathbb{M}\mathcal{A}$ and $a \in \mathcal{A}$, we denote by $d_a : \mathcal{A} \rightarrow \mathbb{R}$ the function “currying” a , defined by $d_a(x) = d(a, x)$. Note that d_a is $d, d_{\mathbb{R}}^{\leq}$ -Lipschitz since $d_a(y) - d_a(x) = d(a, y) - d(a, x) \leq d(x, y)$ follows from the triangle inequality. We are now ready to state the result relating the two distances.

Theorem 1. *Let $d \in \mathbb{M}\mathcal{A}$ be any quasimetric. For all $[a], \alpha \in \mathbb{D}\mathcal{A}$ it holds that*

$$\mathbb{K}^<(d)([a], \alpha) = \mathbb{W}(d)([a], \alpha) = \mathcal{E}_{\alpha} d_a .$$

Proof. We start with the Wasserstein distance. The crucial observation is that for point $[a]$, there is informally a single source “pile of earth”: all the probability has to come from a . As a consequence, $\mathcal{S}_{[a], \alpha}$ contains a unique strategy $S_{x,y} = [a]_x \cdot \alpha_y$ with independent marginals $[a]$ and α . We can then calculate

$$\begin{aligned} & \mathbb{W}(d)([a], \alpha) \\ = & \inf_{S: \mathcal{S}_{[a], \alpha}} \mathcal{E}_S d && \text{“definition of } \mathbb{W} \text{”} \\ = & \mathcal{E}_{(x,y) \sim S} d(x, y) && \text{“take unique } S \text{ with independent marginals } [a], \alpha \text{”} \\ = & \mathcal{E}_{y \sim \alpha} \mathcal{E}_{x \in [a]} d(x, y) && \text{“independence of marginals”} \\ = & \mathcal{E}_{y \sim \alpha} d(a, y) && \text{“expectation over point distribution”} \\ = & \mathcal{E}_{\alpha} d_a \end{aligned}$$

For the Kantorovich distance, we bound it from above by $\mathcal{E}_{\alpha} d_a$, as follows:

$$\begin{aligned}
& \mathbb{K}^<(d)([a], \alpha) \\
= & \sup_{F: \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathcal{A}} \mathcal{E}_{\alpha} F - \mathcal{E}_{[a]} F && \text{“definition of } \mathbb{K}^<\text{”} \\
= & \sup_{F: \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathcal{A}} \mathcal{E}_{\alpha} F - F(a) && \text{“expectation over point distribution”} \\
= & \sup_{F: \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathcal{A}} \mathcal{E}_{x \sim \alpha} (F(x) - F(a)) \\
\leq & \sup_{F: \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathcal{A}} \mathcal{E}_{x \sim \alpha} d(a, x) && \text{“(2), } F \text{ is } d, d_{\mathbb{R}}^{\leq}\text{-Lipschitz”} \\
= & \mathcal{E}_{\alpha} d_a
\end{aligned}$$

Finally we bound $\mathbb{K}^<(d)([a], \alpha)$ from below by $\mathcal{E}_{\alpha} d_a$, as follows:

$$\begin{aligned}
& \mathbb{K}^<(d)([a], \alpha) \\
= & \sup_{F: \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathcal{A}} \mathcal{E}_{\alpha} F - \mathcal{E}_{[a]} F && \text{“definition of } \mathbb{K}^<\text{”} \\
\geq & \mathcal{E}_{\alpha} d_a - \mathcal{E}_{[a]} d_a && \text{“} d_a \in \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathcal{A}\text{”} \\
= & \mathcal{E}_{\alpha} d_a && \text{“} \mathcal{E}_{[a]} d_a = d_a(a) = 0\text{”} \quad \square
\end{aligned}$$

3.2 Computing additive (\mathcal{G}, π) -capacity

We now discuss a generic technique for computing the additive (\mathcal{G}, π) -capacity, for a given family of gain functions $\mathcal{G} \subseteq \mathbb{G}\mathcal{X}$, using the Kantorovich distance. Recall that $\mathcal{ML}_{\mathcal{G}}^+(\pi, C)$ (Def. 1) is defined as the maximum difference between the posterior and prior vulnerabilities $V_g[\pi \triangleright C]$, $V_g[\pi]$. The latter are simply the expected value of V_g over the distributions $[\pi]$, $[\pi \triangleright C]$, which are *hyper* distributions, having sample space $\mathbb{D}\mathcal{X}$.

We start by taking $\mathcal{A} = \mathbb{D}\mathcal{X}$ as the underlying space of the Kantorovich construction. A quasimetric $d \in \mathbb{M}\mathbb{D}\mathcal{X}$ measures the distance between two distributions on secrets. The key for bounding $\mathcal{ML}_{\mathcal{G}}^+(\pi, C)$ from above is to find such a quasimetric d wrt which V_g is Lipschitz for all $g \in \mathcal{G}$. Since the Kantorovich distance maximizes $\mathcal{E}_{\alpha'} F - \mathcal{E}_{\alpha} F$ over all Lipschitz functions F , it will provide an upper bound to the additive capacity.

Bounding the capacity from below is also possible if there exists some $g \in \mathcal{G}$ such that $d_{\pi} = V_g$. This is due to the fact that the g -leakage for this g is exactly $\mathcal{E}_{[\pi \triangleright C]} d_{\pi}$.

In the following, given a class of gain function $\mathcal{G} \subseteq \mathbb{G}\mathcal{X}$, we denote by $V_{\mathcal{G}} = \{V_g \mid g \in \mathcal{G}\}$ the set of g -vulnerabilities induced by \mathcal{G} . The bounding technique is formalized in the following result.

Theorem 2. *Let $d \in \mathbb{M}\mathbb{D}\mathcal{X}$, let $\mathcal{G} \subseteq \mathbb{G}\mathcal{X}$ and fix a channel C and prior π . Then*

$$d_{\pi} \in V_{\mathcal{G}} \quad \text{implies} \quad \mathcal{ML}_{\mathcal{G}}^+(\pi, C) \geq k, \text{ and} \quad (3)$$

$$\mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathbb{D}\mathcal{X} \supseteq V_{\mathcal{G}} \quad \text{implies} \quad \mathcal{ML}_{\mathcal{G}}^+(\pi, C) \leq k, \quad (4)$$

where $k = \mathbb{K}^<(d)([\pi], [\pi \triangleright C]) = \mathbb{W}(d)([\pi], [\pi \triangleright C]) = \mathcal{E}_{[\pi \triangleright C]} d_{\pi}$.

Proof. The fact that $\mathbb{K}^<(d)([\pi], [\pi \triangleright C]) = \mathbb{W}(d)([\pi], [\pi \triangleright C]) = \mathcal{E}_{[\pi \triangleright C]} d_{\pi}$ comes from Thm. 1, for $\mathcal{A} = \mathbb{D}\mathcal{X}$, $a = \pi$, $\alpha = [\pi \triangleright C]$. We start with (3): for $V_g = d_{\pi}$ we have that $V_g(\pi) = 0$ and $V_g[\pi \triangleright C] = \mathcal{E}_{[\pi \triangleright C]} d_{\pi}$ hence $\mathcal{ML}_{\mathcal{G}}^+(\pi, C) \geq \mathcal{L}_{\mathcal{G}}^+(\pi, C) = \mathcal{E}_{[\pi \triangleright C]} d_{\pi}$. For (4) we have that

$$\begin{aligned}
& \mathcal{ML}_{\mathcal{G}}^+(\pi, C) \\
= & \sup_{V_g: V_{\mathcal{G}}} \mathcal{E}_{[\pi \triangleright C]} V_g - \mathcal{E}_{[\pi]} V_g && \text{“definition”} \\
\leq & \sup_{F: \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}} \mathbb{D}\mathcal{X}} \mathcal{E}_{[\pi \triangleright C]} F - \mathcal{E}_{[\pi]} F && \text{“sup over larger class”} \\
= & \mathbb{K}^<(d)([\pi], [\pi \triangleright C]) && \text{“definition”} \quad \square
\end{aligned}$$

So far we have considered an unknown quasimetric d on probability distributions, and identified in Thm. 2 two properties of d that provide bounds for additive leakage. It is not clear, however, whether such a quasimetric exists and what is its relationship with the class \mathcal{G} . We now show that the choice of d is in fact canonical for each class. More precisely, for any \mathcal{G} we can construct a quasimetric $d_{\mathcal{G}}^{\leq}$ satisfying the second condition of Thm. 2. Furthermore, if a quasimetric d satisfying *both* conditions (for any π) does exist, then it is *unique* and equal to $d_{\mathcal{G}}^{\leq}$.

Theorem 3. *Let $\mathcal{G} \subseteq \mathbb{G}\mathcal{X}$ and define a quasimetric $d_{\mathcal{G}}^{\leq}: \mathbb{M}\mathbb{D}\mathcal{X}$ as*

$$d_{\mathcal{G}}^{\leq}(\pi, \sigma) := \sup_{g \in \mathcal{G}} d_{\mathbb{R}}^{\leq}(V_g(\pi), V_g(\sigma)).$$

Then $V_{\mathcal{G}} \subseteq \mathbb{C}^{d_{\mathcal{G}}^{\leq}, d_{\mathbb{R}}^{\leq}}\mathbb{D}\mathcal{X}$. Moreover, if $\{d_{\pi} \mid \pi: \mathbb{D}\mathcal{X}\} \subseteq V_{\mathcal{G}} \subseteq \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}}\mathbb{D}\mathcal{X}$ holds for some quasimetric d , then $d = d_{\mathcal{G}}^{\leq}$.

Proof. Let $g \in \mathcal{G}$. We trivially have that

$$d_{\mathbb{R}}^{\leq}(V_g(\pi), V_g(\sigma)) \leq \sup_{g \in \mathcal{G}} d_{\mathbb{R}}^{\leq}(V_g(\pi), V_g(\sigma)) = d_{\mathcal{G}}^{\leq}(\pi, \sigma),$$

hence V_g is $d_{\mathcal{G}}^{\leq}, d_{\mathbb{R}}^{\leq}$ -Lipschitz. Now let $d: \mathbb{M}\mathbb{D}\mathcal{X}$ such that $\{d_{\pi} \mid \pi: \mathbb{D}\mathcal{X}\} \subseteq V_{\mathcal{G}} \subseteq \mathbb{C}^{d, d_{\mathbb{R}}^{\leq}}\mathbb{D}\mathcal{X}$ and let $\pi, \sigma: \mathbb{D}\mathcal{X}$. From $d_{\pi} \in V_{\mathcal{G}}$, we get that

$$d(\pi, \sigma) = d_{\mathbb{R}}^{\leq}(d_{\pi}(\pi), d_{\pi}(\sigma)) \leq \sup_{g \in \mathcal{G}} d_{\mathbb{R}}^{\leq}(V_g(\pi), V_g(\sigma)) = d_{\mathcal{G}}^{\leq}(\pi, \sigma).$$

Then, since V_g is $d, d_{\mathbb{R}}^{\leq}$ -Lipschitz for all $g \in \mathcal{G}$, we get that

$$d_{\mathcal{G}}^{\leq}(\pi, \sigma) = \sup_{g \in \mathcal{G}} d_{\mathbb{R}}^{\leq}(V_g(\pi), V_g(\sigma)) \leq \sup_{g \in \mathcal{G}} d(\pi, \sigma) = d(\pi, \sigma),$$

hence d and $d_{\mathcal{G}}^{\leq}$ coincide. \square

Finally, an important corollary of this technique is that, assuming that d can be computed in time $O(|\mathcal{X}|)$, $\mathcal{ML}_{\mathcal{G}}^+(\pi, C)$ can be computed in time $O(|\mathcal{X}||\mathcal{Y}|)$. Indeed, calculating $\mathcal{E}_{[\pi \triangleright C]} d_{\pi}$ involves computing the output and posterior distributions of $[\pi \triangleright C]$. The former can be computed in $O(|\mathcal{X}||\mathcal{Y}|)$ time via the joint matrix J ; then for each posterior δ^y , we need to construct δ^y ($O(|\mathcal{X}|)$) and compute $d(\pi, \delta^y)$ ($O(|\mathcal{X}|)$).

3.3 Additive capacity for 1-spanning gain functions

We are now ready to provide a complete method for computing additive capacity for an important family of gain functions. The *span* of a function $f: \mathcal{A} \rightarrow \mathbb{R}$ is defined as

$$\|f\| := \sup_{a, a' \in \mathcal{A}} |f(a) - f(a')|,$$

while for gain functions (having two arguments) we define $\|g\| := \sup_{w \in \mathcal{W}} \|g(w, \cdot)\|$. Since scaling g causes unbounded leakage, a natural solution is to *limit the range of g* . This can be done in an elegant way, without completely fixing g 's range, by requiring

that $\|g\| \leq 1$, which brings us to the class $\mathbb{G}^1\mathcal{X}$ of 1-spanning gain functions, the topic of study in this section. In the following we see that this restriction in fact limits the *steepness* of V_g .

Note that any result for $\mathbb{G}^1\mathcal{X}$ can be straightforwardly extended to k -spanning gain functions, since $\mathcal{L}_{g \times k}^+(\pi, C) = k \cdot \mathcal{L}_g^+(\pi, C)$ implies that the k -spanning additive capacity is simply $k \cdot \mathcal{ML}_{\mathbb{G}^1}^+(\pi, C)$. Note also that this class is quite large: any g with a *finite* number of actions has finite span. For an infinite number of actions, however, it is possible that $\|g\| = +\infty$, i.e. that g is not k -spanning for any k ; important functions such as Shannon vulnerability fall in this category. In §3.4 we enlarge our class of gain functions to include such cases.

Total variation, steepness and g 's span To apply the bounding technique of §3.2 to $\mathbb{G}^1\mathcal{X}$ we need a (quasi)metric $d \in \mathbb{M}\mathbb{D}\mathcal{X}$ with respect to which V_g is Lipschitz when g is 1-spanning. Conveniently, this is the case of the well-known *total variation* distance:

$$\text{tv}(\pi, \pi') := \sup_{X \subseteq \mathcal{X}} |\pi(X) - \pi'(X)|.$$

For discrete distributions, expressed as vectors, the total variation is equal to $d_{\mathbb{R}^n}^{\leq}(\pi, \pi')$, which is in fact symmetric when restricted to probability distributions (because the elements sum up to 1), and equal to $1/2$ the Manhattan distance $\|\pi - \pi'\|_1$. Total variation is a natural choice for $\mathbb{D}\mathcal{X}$ when \mathcal{X} is an “unstructured” space with no underlying metric. Indeed, such spaces can be naturally equipped with the *discrete* metric $\text{dm}: \mathbb{M}\mathcal{X}$, defined as $\text{dm}(x, x') = 0$ iff $x = x'$ and 1 otherwise. It is well known that the Kantorovich lifting of this metric gives total variation, namely $\text{tv} = \mathbb{K}(\text{dm})$. Note, however, that the fact that tv is the *result* of Kantorovich is not important for our goals; our technique involves applying \mathbb{K}^{\leq} to tv itself, lifting it to hyper-distributions.

The Lipschitz property wrt tv and the standard Euclidean $d_{\mathbb{R}}$ naturally expresses the steepness of V_g . If V_g is $k \cdot \text{tv}$ -Lipschitz then the vulnerability can be modified by at most $k \cdot \epsilon$ while changing the probability of any subset of secrets by ϵ . The larger k is, the steeper V_g can be, i.e. the faster it is allowed to change. It turns out that this property is tightly connected to the span of g , as the following result from [1] states.

Proposition 1. *For all $g: \mathbb{G}\mathcal{X}$ it holds that V_g is $\|g\| \cdot \text{tv}$ -Lipschitz.*

As a consequence of the above result we get that $\|V_g\| \leq \|g\|$, since $|V_g(\pi) - V_g(\pi')|$ can be no greater than $\|g\| \cdot \text{tv}(\pi, \pi') \leq \|g\|$.

Putting the pieces together We can finally recover (in a more structured way) the result of [1] for computing the additive (\mathbb{G}^1, π) -capacity. Denote by $\mathbf{1}_S(x)$ the indicator function, equal to 1 if $x \in S$ and 0 otherwise.

Theorem 4. *Given a channel C and prior π , it holds that*

$$\mathcal{ML}_{\mathbb{G}^1}^+(\pi, C) = \mathcal{E}_{[\pi \triangleright C]} \text{tv}_{\pi}.$$

The capacity is realized by the gain function $g: \mathbb{G}^1\mathcal{X}$ defined by

$$\mathcal{W} := 2^{\mathcal{X}}, \quad g(\mathcal{W}, x) := \mathbf{1}_{\mathcal{W}}(x) - \pi(\mathcal{W}),$$

for which it holds that $V_g = \text{tv}_\pi$.⁴

Proof. The result comes from Thm. 2 for $d = \text{tv}$ and $\mathcal{G} = \mathbb{G}^1 \mathcal{X}$; we show here that the two conditions of this theorem hold.

For the upper bound we need to show that $g: \mathbb{G}^1 \mathcal{X}$ implies that V_g is tv , $d_{\mathbb{R}}^{\leq}$ -Lipschitz. But from Prop. 1 we know that V_g is tv , $d_{\mathbb{R}}$ -Lipschitz, which is a *stronger* property since $d_{\mathbb{R}}^{\leq}$ is no greater than $d_{\mathbb{R}}$ for all reals.

For the lower bound, we need to show that the claimed gain function g is 1-spanning and $V_g(\tau) = \text{tv}_\pi(\tau)$. Note that g clearly depends on the fixed π . For the 1-spanning part we have that $|g(W, x) - g(W, x')| = |\mathbf{1}_W(x) - \mathbf{1}_W(x')| \leq 1$. Moreover, it holds that

$$\begin{aligned}
 & V_g(\tau) \\
 = & \max_W \mathcal{E}_{x \sim \tau} g(w, x) && \text{“definition of } V_g \text{”} \\
 = & \max_W (\mathcal{E}_{x \sim \tau} \mathbf{1}_W(x) - \pi(W)) && \text{“definition of } g \text{”} \\
 = & \max_W \sum_{x \in W} (\tau_x - \pi_x) \\
 = & \sum_x d_{\mathbb{R}}^{\leq}(\pi_x, \tau_x) && \text{“take } W = \{x \mid \tau_x \geq \pi_x\} \text{”} \\
 = & \text{tv}(\pi, \tau) . && \text{“}\text{tv}(\pi, \pi') = d_{\mathbb{R}^n}^{\leq}(\pi, \pi') \text{”} \quad \square
 \end{aligned}$$

In the above proof we showed that tv satisfies the two conditions of Thm. 2. From Thm. 3 we know that there is a unique quasimetric satisfying these conditions which can be constructed explicitly from the class $\mathcal{G} = \mathbb{G}^1 \mathcal{X}$, that is: $\text{tv}(\pi, \pi') = d_{\mathbb{G}^1}^{\leq}(\pi, \sigma) = \sup_{g: \mathbb{G}^1 \mathcal{X}} V_g(\sigma) - V_g(\pi)$. Note also that tv can be computed in $|\mathcal{X}|$ time, hence, as discussed in §3.2, $\mathcal{ML}_{\mathbb{G}^1}^+(\pi, C)$ can be computed in time $O(|\mathcal{X}| |\mathcal{Y}|)$.

3.4 Additive capacity for 1-spanning vulnerability functions

As discussed in the introduction it is often desirable to measure vulnerability within a predefined range, for instance $[0, 1]$ or $[0, \log_2 n]$ ($n = |\mathcal{X}|$). A natural way to achieve this is to consider k -spanning gain functions, implicitly limiting the range of V_g to an interval of size at most k . This choice, however, excludes important vulnerabilities that cannot be expressed as V_g for any k -spanning g .

For instance, for the Shannon vulnerability function V_H (see §2) the additive capacity is equal to the well known Shannon mutual information. Although V_H lies within $[0, \log_2 n]$ and it can be expressed as V_{g_H} for a suitable g_H , this gain function is *not* $\log_2 n$ -spanning, in fact $\|g_H\| = +\infty$. As a consequence, the additive capacity $\mathcal{ML}_{\mathbb{G}^1}^+(\pi, C)$, discussed in the previous section, does not provide a bound for g_H -leakage. Indeed, the mutual-information of the fully transparent identity channel C_{id} on a uniform prior is $\mathcal{L}_{g_H}^+(\pi^u, C_{\text{id}}) = \log_2 n$, which exceeds its additive capacity wrt $\log_2 n$ -spanning gain functions, which is equal to $\log_2 n \cdot \mathcal{ML}_{\mathbb{G}^1}^+(\pi^u, C_{\text{id}}) = \frac{n-1}{n} \log_2 n$.

Aiming at robustness wrt a larger class of vulnerabilities, we can allow functions V_g that have a bounded range, even though g itself is unbounded. Similarly to $\mathbb{G}^1 \mathcal{X}$, we choose to *limit* the range of V_g without completely fixing it, by restricting to the class $\mathbb{G}^\dagger \mathcal{X} = \{g: \mathbb{G} \mathcal{X} \mid \|V_g\| \leq 1\}$ of *1-spanning vulnerability functions*.

⁴ The choice $\mathcal{W} = \mathcal{X} \rightarrow \{-1, 1\}$, $g(w, x) = \frac{1}{2}(w(x) - \mathcal{E}_\pi w)$ is also capacity-realizing [1].




Fig. 2. Line connecting π and σ , extended to the boundary of the simplex.

Since $\|V_g\| \leq \|g\|$, but not vice-versa (as V_H demonstrates), it holds that $\mathbb{G}^1\mathcal{X} \subset \mathbb{G}^\dagger\mathcal{X}$. Since any convex (and continuous) function can be expressed as V_g for some properly constructed g [2], $V_{\mathbb{G}^\dagger\mathcal{X}}$ is the class of all 1-spanning convex functions. Note also that, since any $V_g, g: \mathbb{G}\mathcal{X}$ is bounded, it is k -spanning for some k .

To compute the additive capacity via the technique of §3.2, we need a quasimetric satisfying both conditions of Thm. 2. From Theorem 3 we know that such a quasimetric (if it exists) is unique and equal to the $d_{\mathbb{G}}^<$ construction. In the previous section, this turned out to be the well-known total variation distance. In this section, on the other hand, we start directly with $d_{\mathbb{G}}^<$ for our class $\mathcal{G} = \mathbb{G}^\dagger\mathcal{X}$. The resulting quasimetric $d_{\mathbb{G}^\dagger}^<$ is called the ‘‘convex-separation’’ quasimetric, and is given by

$$d_{\mathbb{G}^\dagger}^<(\pi, \sigma) := \sup_{g \in \mathbb{G}^\dagger\mathcal{X}} d_{\mathbb{R}}^<(V_g(\pi), V_g(\sigma)) = \sup_{g \in \mathbb{G}^\dagger\mathcal{X}} V_g(\sigma) - V_g(\pi).$$

Note that, once again, we removed the max from the definition of $d_{\mathbb{R}}^<$ since the sup is anyway non-negative.

An important property of $d_{\mathbb{G}^\dagger}^<$ is that it admits a simple closed-form solution.

Theorem 5. *The convex-separation quasimetric $d_{\mathbb{G}^\dagger}^<$ is equal to*

$$d_{\mathbb{G}^\dagger}^<(\pi, \sigma) = \max_{x: \lceil \pi \rceil} 1 - \frac{\sigma_x}{\pi_x}.$$

Proof. Let $\pi, \sigma \in \mathbb{D}\mathcal{X}$. Assume $\pi \neq \sigma$ (the case $\pi = \sigma$ is trivial) and consider the line in \mathbb{R}^n joining the two priors, as shown in Fig. 2. The points on that line, starting from π and moving towards σ can be written as $\pi^c = \pi + c(\sigma - \pi)$ for $c \geq 0$. Continuing on that line, at some point we are going to hit the boundary of the probability simplex $\mathbb{D}\mathcal{X}$. Let π^c be the point on that boundary, i.e.

$$\mathbf{c} := \max\{c \mid \pi^c \in \mathbb{D}\mathcal{X}\}. \quad (5)$$

Note that $\mathbf{c} \geq 1$ since $\pi^1 = \sigma \in \mathbb{D}\mathcal{X}$. Now let $F: \mathbb{D}\mathcal{X} \rightarrow \mathbb{R}$ be convex and 1-spanning. Since σ lies in the line segment between π and π^c , we can write it as a convex combination

$$\sigma = \mathbf{c}^{-1}\pi^c + (1 - \mathbf{c}^{-1})\pi \quad (6)$$

with $\mathbf{c}^{-1} \in (0, 1]$. From convexity we get that

$$F(\sigma) \leq \mathbf{c}^{-1}F(\pi^c) + (1 - \mathbf{c}^{-1})F(\pi),$$

from which, together with $F(\pi^c) - F(\pi) \leq 1$ (F is 1-spanning), we get

$$F(\sigma) - F(\pi) \leq \mathbf{c}^{-1}(F(\pi^c) - F(\pi)) \leq \mathbf{c}^{-1}. \quad (7)$$

We now compute \mathbf{c} , which is given by the maximization problem (5). The problem has a single variable c and the constraint $\pi^c \in \mathbb{D}\mathcal{X}$ can be expressed by $\sum_x \pi_x^c = 1$ and the inequalities $\pi_x^c \geq 0$. The first constraint $\sum_x \pi_x^c = 1$ is always satisfied by construction of π^c . Hence we only need to ensure that $\pi_x^c = \pi_x + c(\sigma_x - \pi_x) \geq 0$ for all $x: \mathcal{X}$. If $\pi_x = \sigma_x$ this is always satisfied, and if $\pi_x < \sigma_x$ then this imposes a *lower* bound on c . The only interesting case is when $\pi_x > \sigma_x$ which gives us an upper bound $c \leq \pi_x / (\pi_x - \sigma_x)$. The max c satisfying all upper bounds is equal to the smallest of them:

$$\mathbf{c} = \min_{x: \pi_x > \sigma_x} \frac{\pi_x}{\pi_x - \sigma_x}.$$

Replacing \mathbf{c} in (7) we get

$$F(\sigma) - F(\pi) \leq \max_{x: \pi_x > \sigma_x} \frac{\pi_x - \sigma_x}{\pi_x} = \max_{x: \lceil \pi \rceil} 1 - \frac{\sigma_x}{\pi_x}.$$

We finally show that the above bound is attainable. Define

$$F_\pi(\tau) := \max_{x: \lceil \pi \rceil} 1 - \frac{\tau_x}{\pi_x}.$$

F_π is convex as the max of convex (in fact linear) functions of τ , so it can be expressed as V_g (see Thm. 6 for the exact g). Moreover, $F_\pi(\pi) = 0$, hence $F_\pi(\sigma) - F_\pi(\pi) = \max_{x: \lceil \pi \rceil} 1 - \frac{\sigma_x}{\pi_x}$, which concludes the proof. \square

We can now use $d_{\mathbb{G}^\dagger}^<$ to compute the additive $(\mathbb{G}^\dagger, \pi)$ -capacity.

Theorem 6. *Given a channel C and prior π , it holds that*

$$\mathcal{ML}_{\mathbb{G}^\dagger}^+(\pi, C) = \mathcal{E}_{[\pi \triangleright C]} d_{\mathbb{G}^\dagger}^<_{\pi} = 1 - \sum_{y: \mathcal{Y}} \min_{x: \lceil \pi \rceil} C_{x,y}.$$

The capacity is realized by the complement of the “ π -reciprocal” gain function

$$\mathcal{W} = \lceil \pi \rceil, \quad g_{\pi^{-1}}^c = \begin{cases} 1 - \frac{1}{\pi_x}, & \text{if } w = x \\ 1, & \text{if } w \neq x \end{cases},$$

for which it holds that $V_{g_{\pi^{-1}}^c} = d_{\mathbb{G}^\dagger}^<_{\pi}$, and as a consequence $g_{\pi^{-1}}^c: \mathbb{G}^\dagger \mathcal{X}$.

Proof. The result comes from Thm. 2 for $d = d_{\mathbb{G}^\dagger}^<$ and $\mathcal{G} = \mathbb{G}^\dagger \mathcal{X}$; we show here that its two conditions of the theorem hold. The second is satisfied automatically by the construction of $d_{\mathbb{G}^\dagger}^<$ (Thm. 3). For the first condition, after simple calculations we find that the $g_{\pi^{-1}}^c$ -vulnerability function is equal to $V_{g_{\pi^{-1}}^c}(\sigma) = \max_{x: \lceil \pi \rceil} 1 - \frac{\sigma_x}{\pi_x}$, hence from Thm. 5 we have that $V_{g_{\pi^{-1}}^c} = d_{\mathbb{G}^\dagger}^<_{\pi}$. Finally, simple calculations show that $V_{g_{\pi^{-1}}^c}[\pi \triangleright C] = 1 - \sum_y \min_{x \in \lceil \pi \rceil} C_{x,y}$, which concludes the proof since $V_{g_{\pi^{-1}}^c}(\pi) = d_{\mathbb{G}^\dagger}^<_{\pi}(\pi) = 0$. \square

Note that the capacity-realizing gain function $g_{\pi^{-1}}^c$ essentially “cancels out” the effect of the prior, making $\mathcal{ML}_{\mathbb{G}^\dagger}^+(\pi, C)$ independent from π , and equal to 1 minus the sum of the column *minima* of C (ignoring the rows when $\pi_x = 0$). Remarkably,

the “ π -reciprocal” gain function $g_{\pi^{-1}}$ (the complement of $g_{\pi^{-1}}^c$) produces the same “cancellation” effect for *multiplicative* (\mathbb{G}^+, π) -capacity, making it independent from π .

An observation that can be made from Thm. 6 is that the capacity realizing g is k -spanning for $k = \max_{x: \lceil \pi \rceil} \frac{1}{\pi_x}$. From this we can conclude that $\mathcal{ML}_{\mathbb{G}^\dagger}^+(\pi, C) \leq k \cdot \mathcal{ML}_{\mathbb{G}^1}^+(\pi, C)$ for $k = \max_{x: \lceil \pi \rceil} 1/\pi_x$. In particular $\mathcal{ML}_{\mathbb{G}^\dagger}^+(\pi^u, C) \leq |\mathcal{X}| \cdot \mathcal{ML}_{\mathbb{G}^1}^+(\pi^u, C)$ for the uniform π^u .

A final note about our use of quasimetrics. Although the total variation tv , used in §3.3 for $\mathbb{G}^1\mathcal{X}$, is a proper metric, $d_{\mathbb{G}^\dagger}^<$ used in this section for $\mathbb{G}^\dagger\mathcal{X}$ is not, since it is not symmetric. This is why we had to work with quasimetrics; it is certainly possible to define a symmetric variant of $d_{\mathbb{G}^\dagger}^<$ (eg. as $\max\{d_{\mathbb{G}^\dagger}^<(\pi, \sigma), d_{\mathbb{G}^\dagger}^<(\sigma, \pi)\}$), however this metric would not satisfy both conditions of Thm. 2. Recall that for each class \mathcal{G} there can be at most one quasimetric satisfying both properties, and for $\mathbb{G}^\dagger\mathcal{X}$ this is $d_{\mathbb{G}^\dagger}^<$.

3.5 Maximize over both g and π

This scenario was left open in [1] (which uses the class $\mathbb{G}^1\mathcal{X}$), since $\mathcal{ML}_{\mathbb{G}^1}^+(\pi, C)$ depends on the prior, and maximizing it over π is challenging. Our results on the larger class $\mathbb{G}^\dagger\mathcal{X}$, however, lead to a complete solution since $\mathcal{ML}_{\mathbb{G}^\dagger}^+(\pi, C)$ is independent from π . By Thm. 6, any full-support π with $g_{\pi^{-1}}^c$ are capacity-realizing, giving

$$\mathcal{ML}_{\mathbb{G}^\dagger}^+(\mathbb{D}, C) = 1 - \sum_y \min_x C_{x,y}.$$

4 The additive miracle theorem

The multiplicative Bayes-capacity $\mathcal{ML}_{g_{\text{id}}}^\times(\mathbb{D}, C)$ is well known to be realized on a uniform prior, and is equal to the sum of the column maxima of C [9, 24]. A result from [3], which was surprising enough to be named “Miracle”, states that $\mathcal{ML}_{g_{\text{id}}}^\times(\mathbb{D}, C)$ is in fact a *universal upper bound* for multiplicative leakage (wrt non-negative g 's).

Theorem 7 (“Miracle”, [3]). *For any C , $\pi: \mathbb{D}\mathcal{X}$, and non-negative $g: \mathbb{G}^+\mathcal{X}$, we have*

$$\mathcal{L}_g^\times(\pi, C) \leq \sum_y \max_x C_{x,y} = \mathcal{ML}_{g_{\text{id}}}^\times(\mathbb{D}, C).$$

In [1], this theorem was used to easily conclude that the capacity $\mathcal{ML}_{\mathbb{G}^+}^\times(\pi, C)$ is equal to $\mathcal{ML}_{g_{\text{id}}}^\times(\mathbb{D}, C)$ for any full-support π . In the additive case, having already a solution for $\mathcal{ML}_{\mathbb{G}^\dagger}^+(\pi, C)$, we can go in the opposite direction and obtain an additive variant of the miracle theorem. Denote by $g_{\text{id}}^c = 1 - g_{\text{id}}$ the complement of g_{id} .

Theorem 8 (“Additive Miracle”). *For any C , $\pi: \mathbb{D}\mathcal{X}$, and $g: \mathbb{G}^\dagger\mathcal{X}$, we have*

$$\mathcal{L}_g^+(\pi, C) \leq 1 - \sum_y \min_x C_{x,y} = |\mathcal{X}| \cdot \mathcal{ML}_{g_{\text{id}}^c}^+(\mathbb{D}, C).$$

Proof. The inequality is a direct consequence of Thm. 6; note that it holds for any prior since $1 - \sum_{y: \mathcal{Y}} \min_{x: \mathcal{X}} C_{x,y} \geq 1 - \sum_{y: \mathcal{Y}} \min_{x: \lceil \pi \rceil} C_{x,y}$. Now let $g^* = g_{\text{id}}^c \times |\mathcal{X}|$, for which it holds that $\mathcal{ML}_{g^*}^+(\mathbb{D}, C) = |\mathcal{X}| \cdot \mathcal{ML}_{g_{\text{id}}^c}^+(\mathbb{D}, C)$. We have that $V_{g^*} = |\mathcal{X}|(1 - \min_{x: \mathcal{X}} \pi_x)$. For uniform π^u we compute $\mathcal{L}_{g^*}^+(\pi^u, C) = 1 - \sum_y \min_x C_{x,y}$, and since $g^*: \mathbb{G}^\dagger\mathcal{X}$, this is an upper bound for all $\pi: \mathbb{D}\mathcal{X}$, and hence equal to $\mathcal{ML}_{g^*}^+(\mathbb{D}, C)$. \square

The multiplicative and additive miracle theorems are similar in nature, although they do have several differences. They both provide a universal bound for leakage, which holds for all priors and all gain functions within a certain class. In the multiplicative case, this is the class $\mathbb{G}^+\mathcal{X}$ of non-negative gain functions, while in the additive case, the class $\mathbb{G}^\dagger\mathcal{X}$ of gain functions producing a 1-spanning V_g . In the multiplicative case the bound is given by the sum of column maxima of C , while in the additive case by 1 minus the sum of column minima. In the multiplicative case the bound coincides with the $(g_{\text{id}}, \mathbb{D})$ -capacity for the identity gain function (i.e. the Bayes-capacity), which is realized on a uniform prior. In the additive case the bound is $(|\mathcal{X}| \text{ times})$ the $(g_{\text{id}}^c, \mathbb{D})$ -capacity for the “complement of identity” gain function, also realized on a uniform prior.

Example Consider the case $\mathcal{X} = \{x_1, x_2\}$ with a gain function penalizing wrong guesses, defined as $\mathcal{W} = \mathcal{X}$, and $g(w, x) = 1$ iff $w = x$ and -1 otherwise. Note that V_g is always non-negative since the probability of a correct guess is at least 0.5 (for $|\mathcal{X}| = 2$). For a uniform prior π^u , both guesses are equivalent, giving expected gain $0.5 \cdot 1 + 0.5 \cdot -1 = 0$, so $V_g(\pi^u) = 0$.

C	y_1	y_2
x_1	0.8	0.2
x_2	0.2	0.8

Now consider the illustrated channel C which gives rather good information about the secret. Computing the two posteriors we get $\delta^{y_1} = (0.8, 0.2)$ and $\delta^{y_2} = (0.2, 0.8)$, which both give g -vulnerability $V_g(\delta^{y_1}) = V_g(\delta^{y_2}) = 0.8 - 0.2 = 0.6$. Hence $V_g[\pi^u \triangleright C] = 0.6$ and as a consequence $\mathcal{L}_g^\times(\pi^u, C) = +\infty$, clearly larger than the multiplicative Bayes capacity $\mathcal{ML}_{g_{\text{id}}}^\times(\mathbb{D}, C) = 0.8 + 0.8 = 1.6$. The miracle theorem does not apply here since g takes negative values.

On the other hand, V_g is 1-spanning (although g itself is 2-spanning), since its value is at least 0 (for a uniform prior) and at most 1 (for a point prior). As a consequence the additive miracle theorem applies, guaranteeing that $\mathcal{L}_g^+(\pi^u, C) \leq 1 - \sum_y \min_x C_{x,y} = 1 - 0.2 - 0.2 = 0.6$. Indeed $\mathcal{L}_g^+(\pi^u, C) = 0.6 - 0$, exactly matching the bound. \square

5 Conclusion and future work

We studied the problem of computing additive g -capacities. Extending the Kantorovich technique of [1] with quasimetrics, we provided a solution for the class $\mathbb{G}^\dagger\mathcal{X}$ of 1-spanning vulnerabilities, which, in contrast to $\mathbb{G}^1\mathcal{X}$, can include any vulnerability function (by scaling). The results also provided a solution to the problem of maximizing leakage over both π and g , and lead to an additive variant of the miracle theorem of [3].

In future work we plan to study approximation algorithms for all scenarios, especially $\mathcal{ML}_g^+(\mathbb{D}, C)$ which is NP-complete in general. Moreover, we aim at developing a theory that unifies the two main approaches to robustness, namely *capacity* and *refinement*.

References

1. Alvim, M.S., Chatzikokolakis, K., McIver, A., Morgan, C., Palamidessi, C., Smith, G.: Additive and multiplicative notions of leakage, and their capacities. In: Proc. of CSF. pp. 308–322. IEEE (2014)
2. Alvim, M.S., Chatzikokolakis, K., McIver, A., Morgan, C., Palamidessi, C., Smith, G.: Axioms for information leakage. In: Proc. of CSF. pp. 77–92 (2016)

3. Alvim, M.S., Chatzikokolakis, K., Palamidessi, C., Smith, G.: Measuring information leakage using generalized gain functions. In: Proc. of CSF. pp. 265–279 (2012)
4. Antonopoulos, T., Gazzillo, P., Hicks, M., Koskinen, E., Terauchi, T., Wei, S.: Decomposition instead of self-composition for proving the absence of timing channels. In: PLDI. pp. 362–375. ACM (2017)
5. Backes, M., Köpf, B., Rybalchenko, A.: Automatic discovery and quantification of information leaks. In: Proc. of S&P. pp. 141–153 (2009)
6. Barthe, G., Köpf, B.: Information-theoretic bounds for differentially private mechanisms. In: Proc. of CSF. pp. 191–204 (2011)
7. Biondi, F., Kawamoto, Y., Legay, A., Traonouez, L.: Hyleak: Hybrid analysis tool for information leakage. In: ATVA. LNCS, vol. 10482, pp. 156–163. Springer (2017)
8. Biondi, F., Legay, A., Traonouez, L.M., Wąsowski, A.: Quail: A quantitative security analyzer for imperative code. In: Proc. of CAV. pp. 702–707. Springer (2013)
9. Braun, C., Chatzikokolakis, K., Palamidessi, C.: Quantitative notions of leakage for one-try attacks. In: Proc. of MFPS. ENTCS, vol. 249, pp. 75–91. Elsevier (2009)
10. Chatzikokolakis, K., Chothia, T., Guha, A.: Statistical measurement of information leakage. In: Proc. of TACAS. pp. 390–404 (2010)
11. Chatzikokolakis, K., Palamidessi, C., Panangaden, P.: On the Bayes risk in information-hiding protocols. *J. of Comp. Security* **16**(5), 531–571 (2008)
12. Clarkson, M.R., Schneider, F.B.: Quantification of integrity. In: Proc. of CSF. pp. 28–43 (2010)
13. Cover, T.M., Thomas, J.A.: *Elements of Information Theory*. J. Wiley & Sons, Inc., second edn. (2006)
14. Doychev, G., Köpf, B.: Rigorous analysis of software countermeasures against cache attacks. In: PLDI. pp. 406–421. ACM (2017)
15. Heusser, J., Malacaria, P.: Quantifying information leaks in software. In: Proc. ACSAC '10. pp. 261–269 (2010)
16. Khouzani, M.H.R., Malacaria, P.: Relative perfect secrecy: Universally optimal strategies and channel design. In: Proc. of CSF. pp. 61–76 (2016)
17. Köpf, B., Basin, D.: An information-theoretic model for adaptive side-channel attacks. In: Proc. of CCS. pp. 286–296 (2007)
18. Köpf, B., Mauborgne, L., Ochoa, M.: Automatic quantification of cache side-channels. In: Proc. of CAV '12. pp. 564–580 (2012)
19. Köpf, B., Rybalchenko, A.: Approximation and randomization for quantitative information-flow analysis. In: Proc. of CSF. pp. 3–14 (2010)
20. Köpf, B., Smith, G.: Vulnerability bounds and leakage resilience of blinded cryptography under timing attacks. In: Proc. of CSF. pp. 44–56 (2010)
21. Malacaria, P.: Assessing security threats of looping constructs. In: Proc. of POPL. pp. 225–235 (2007)
22. McIver, A., Meinicke, L., Morgan, C.: Compositional closure for Bayes risk in probabilistic noninterference. In: Proc. ICALP'10. pp. 223–235 (2010)
23. Meng, Z., Smith, G.: Calculating bounds on information leakage using two-bit patterns. In: Proc. of PLAS. pp. 1:1–1:12 (2011)
24. Smith, G.: On the foundations of quantitative information flow. In: Proc. of FOSSACS. LNCS, vol. 5504, pp. 288–302. Springer (2009)
25. Sweet, I., Trilla, J.M.C., Scherrer, C., Hicks, M., Magill, S.: What's the over/under? probabilistic bounds on information leakage. In: POST. LNCS, vol. 10804, pp. 3–27. Springer (2018)
26. Villani, C.: *Topics in optimal transportation*. No. 58, American Mathematical Soc. (2003)
27. Yasuoka, H., Terauchi, T.: Quantitative information flow — verification hardness and possibilities. In: Proc. of CSF. pp. 15–27 (2010)