Axiomatics

Mário S. Alvim
(msalvim@dcc.ufmg.br)

Quantitative Information Flow

DCC-UFMG
(2021/1)

Introduction

- In previous chapters we discussed how secrecy can be quantified:
  - we described information measures that map a prior to a real number reflecting the amount of threat to which the secret is subjected; and
  - we introduced g-vulnerabilities as a rich family of such information measures that can capture a variety of significant operational scenarios.

  In particular, we determined that:
  - g-vulnerabilities can express Bayes vulnerability as a special case, and
  - by using properly constructed loss functions, we can express Shannon entropy and guessing entropy.
- That suggests that g-vulnerabilities constitute a quite general family of information measures.

  But we haven't yet developed independently a clear sense of what a general definition of "vulnerability measure" ought to be.

Introduction

- Here we investigate two significant questions regarding how general g-vulnerabilities really are, as information measures.

  1. Are all instances of g-vulnerabilities "reasonable"?
  2. Are there "reasonable" vulnerability measures that cannot be captured as g-vulnerabilities?

  We address those questions by considering a set of axioms that characterize intuitively reasonable properties that vulnerability measures might satisfy.

  We separately consider:

  - axioms for prior vulnerability,
  - axioms for posterior vulnerability, and
  - axioms for the relationship between prior and posterior vulnerability.

- As a result, we are able to derive properties of leakage.

An axiomatic view of vulnerability

An axiomatic view of vulnerability

- Earlier in this course, we
  - derived vulnerability measures by quantifying the adversary's success in specific operational scenarios; and
  - extended vulnerability measures on priors to vulnerability measures on posteriors, by averaging vulnerabilities over the hypers.
- Now we'll follow a different approach: we axiomatize the study of the vulnerabilities themselves.
- We begin by considering generic vulnerability functions of type

$$\textit{prior vulnerability} — \quad \mathbb{V}\colon \mathbb{D}\mathcal{X} \to \mathbb{R}_{\geq 0}$$

$$\textit{posterior vulnerability} — \quad \widehat{\mathbb{V}}\colon \mathbb{D}^2\mathcal{X} \to \mathbb{R}_{\geq 0}$$

and we examine a variety of properties that "reasonable" instantiations of those generic functions might be expected to have.

We then formalize those properties as a set of **axioms** for vulnerability functions, and investigate their consequences.

An axiomatic view of vulnerability

- We clarify that we are not treating axiomatics here as "self-evident truths".
- A variety of (sets of) axioms might appear intuitively reasonable.

  While it is sensible to consider justifications for each of them, such justifications should not be considered absolute.
- Rather, the axiomatics help us better to understand the logical dependencies among different properties, so that we might identify a minimal set of axioms sufficient to imply all the properties we care about.

An axiomatic view of vulnerability

- The axioms we'll consider are summarized as follows.

  (Concrete channels are used when explicit matrix operations are needed.)

Axioms for prior vulnerabilities $\mathbb{V}$		
CNTY	$\forall \pi$:	$\mathbb{V}$ is a continuous function of π
CVX	$\forall \sum_i a_i \pi^i$:	$\mathbb{V}\left(\sum_i a_i \pi^i\right) \leq \sum_i a_i \mathbb{V}(\pi^i)$
Q-CVX	$\forall \sum_i a_i \pi^i$:	$\mathbb{V}\left(\sum_i a_i \pi^i\right) \leq \max_i \mathbb{V}(\pi^i)$

Axioms for posterior vulnerabilities $\widehat{\mathbb{V}}$		
NI	$\forall \pi$:	$\widehat{\mathbb{V}}[\pi] = \mathbb{V}(\pi)$
DPI	$\forall \pi, \mathsf{C}, \mathsf{R}$:	$\widehat{\mathbb{V}}[\pi \triangleright \mathsf{C}] \geq \widehat{\mathbb{V}}[\pi \triangleright \mathsf{CR}]$
MONO	$\forall \pi, C$:	$\widehat{\mathbb{V}}[\pi \triangleright C] \geq \mathbb{V}(\pi)$

Possible relationships between $\mathbb{V}$ and $\widehat{\mathbb{V}}$		
AVG	$\forall \Delta$:	$\widehat{\mathbb{V}}\Delta = \mathcal{E}_\Delta \mathbb{V}$
MAX	$\forall \Delta$:	$\widehat{\mathbb{V}}\Delta = \max_{\lceil \Delta \rceil} \mathbb{V}$

Axiomatization of prior vulnerabilities

Axiomatization of prior vulnerabilities

- We begin by introducing axioms that deal solely with generic prior vulnerabilities $\mathbb{V}$.
- The first property we consider is that "small" changes on the prior π have a "small" effect on $\mathbb{V}$ applied to that prior.

  That intuition is formalized in the following axiom.

Definition (11.1 - Axiom of continuity (CNTY))

A vulnerability $\mathbb{V}$ is a continuous function of π with respect to the standard topology on $\mathbb{D}\mathcal{X}$, where by *standard* topology on $\mathbb{D}\mathcal{X}$ we mean that it is generated by the *Manhattan metric*

$$d(\pi, \pi') \;:=\; \sum_{x:\,\mathcal{X}} |\pi_x - \pi'_x| \quad .$$

- Intuitively, the CNTY axiom describes adversaries who are not infinitely risk-averse.

Axiomatization of prior vulnerabilities

- Example 1 For instance, the non-continuous vulnerability function

$$\mathbb{V}^{\lambda}(\pi) \;:=\; \begin{cases} 1, & \text{if } \max_x \pi_x \geq \lambda \\ 0, & \text{otherwise} \end{cases}$$

would correspond to an adversary who requires the probability of guessing correctly to be above a certain threshold λ in order to consider an attack effective at all.

But this is arguably an unnatural behavior if we assume that the risk to the adversary of changing the probability to $\lambda - \epsilon$, for negligible ϵ, should not be arbitrarily large. ◁

Axiomatization of prior vulnerabilities

- The second property we consider is that $\mathbb{V}$ is a convex function of the prior.

  (Recall that a convex combination of distributions is itself a distribution.)

Definition (11.2 - Axiom of convexity (CVX))

A vulnerability $\mathbb{V}$ is a convex function of π — that is, for all convex combinations $\sum_i a_i \pi^i$ of distributions we have

$$\mathbb{V}\left(\sum_n a_n \pi^n\right) \quad \leq \quad \sum_n a_n \mathbb{V}(\pi^n) \quad .$$

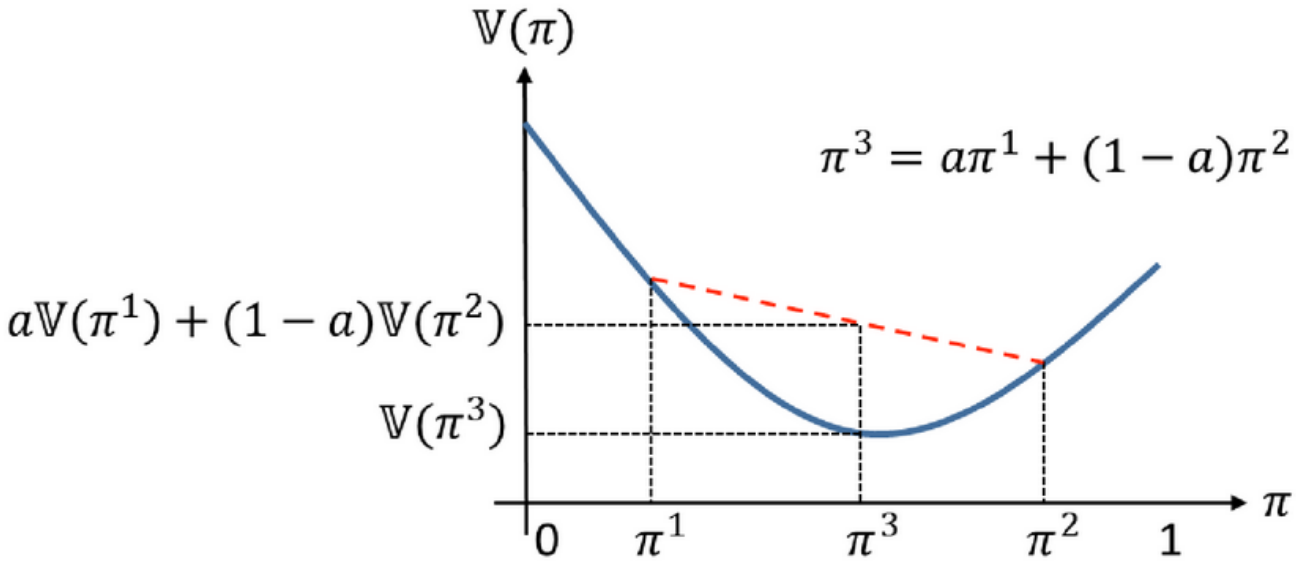

Axiomatization of prior vulnerabilities

- The CVX axiom can be motivated as follows.
- Consider a game in which a secret X in $\mathcal{X}$ (say a password) is drawn from one of two possible prior distributions π^1 and π^2 in $\mathbb{D}\mathcal{X}$.

  More precisely, the choice of prior distribution is itself random:

  1. we first select n in $\{1, 2\}$ with n=1 having probability say a_1 and n=2 probability $a_2 = 1-a_1$, and
  2. after that, we use π^n to draw the actual secret x.

  Now consider the following two scenarios for this game:

  - First scenario: the value of n (i.e. the selection of prior) is revealed to the adversary.

    Using information in that π^n (whichever one it was) the adversary performs an attack on X whose expected success is therefore quantified as $\mathbb{V}(\pi^n)$.

    The expected measure of success overall will therefore be

$$a_1\mathbb{V}(\pi^1) + a_2\mathbb{V}(\pi^2) \ .$$

Axiomatization of prior vulnerabilities

- Now consider the following two scenarios for this game:
  - Second scenario: the choice n is not disclosed to the adversary; instead she knows only that, on average, the secret is drawn from the "overall" prior $a_1\pi^1 + a_2\pi^2$.

    With only that knowledge, her expected success is now quantified as

    $$\mathbb{V}(a_1\pi^1 + a_2\pi^2) \ .$$

  In this game, the CVX axiom corresponds to the intuition that, since in the first scenario the adversary has more information, the effectiveness of an attack can never be smaller.

- A second way of understanding this axiom is to realize that an adversary should get no less information from a_1, π^1 and π^2 than from $a_1\pi^1 + a_2\pi^2$, since the last value can be calculated if the first three are known.

- Note that, in the definition of CVX, it is sufficient to use convex combinations of just two priors, i.e. of the form $a\pi^1 + (1-a)\pi^2$ as in our example above (where a_2 was a_1-1); and indeed we often use such combinations in proofs.

Axiomatization of prior vulnerabilities

- Note that since the vulnerabilities $\mathbb{V}(\pi^n)$ in CVX are weighted by the probabilities a_n, we could have cases when the expected vulnerability $\sum_n a_n \mathbb{V}(\pi^n)$ is small even though some individual $\mathbb{V}(\pi^n)$ is large.

  Then one might argue that the bound imposed by CVX is too strict.

- One could then offer an alternative, more relaxed, axiom, requiring only that $\mathbb{V}\left(\sum_i a_i \pi^i\right)$ is bounded only by the maximum of the individual vulnerabilities.

  That weaker requirement is formalized as the following axiom.

Definition (11.3 - Axiom of quasiconvexity (Q-CVX))

A vulnerability $\mathbb{V}$ is a quasiconvex function of π — i.e. for all convex combinations $\sum_n a_n \pi^n$ we have

$$\mathbb{V}\left(\sum_n a_n \pi^n\right) \quad \leq \quad \max_n \mathbb{V}(\pi^n) \quad .$$

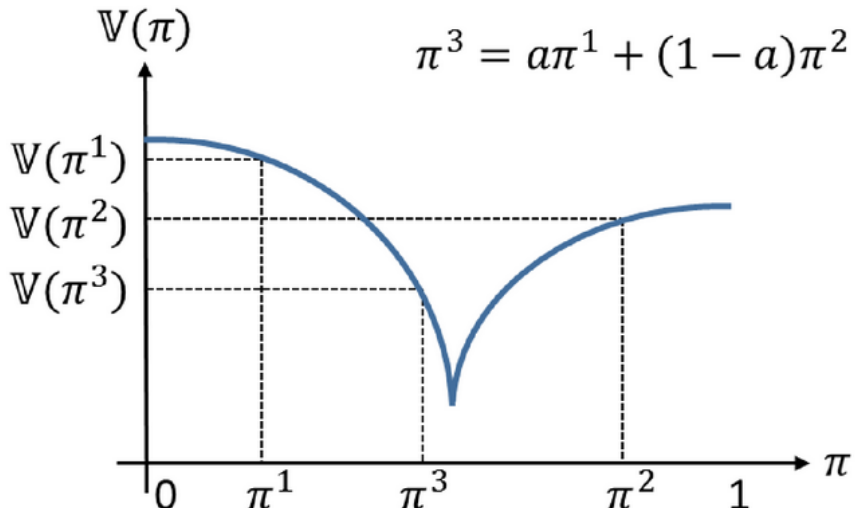

Axiomatization of prior vulnerabilities: Soundness and completeness of V_g wrt. continuous, convex functions

- We will soon discuss further the justification of CVX and Q-CVX as two possible generic properties for studying information flow.
- But before that we provide some results concerning the characterization of g-vulnerabilities in terms of convexity and continuity.
- In particular we show that any real-valued function over priors that is both convex and continuous can be expressed as a g-vulnerability for some g, provided we allow infinitely many possible choices $\mathcal{W}$.

  That implies that g-vulnerabilities are fundamental to an understanding of information flow when CVX is taken as axiomatic.

Axiomatization of prior vulnerabilities: Soundness and completeness of V_g wrt. continuous, convex functions

- The characterization of convex and continuous functions as g-vulnerabilities can be graphed.
  - Each action available to an adversary is a line.

    I.e., a piecewise linear part of V_g as a real-valued function of π.
  - More generally, a vulnerability V_g over an arbitrary $\mathbb{D}\mathcal{X}$ satisfies `CVX`, but it's also continuous if it's finite over its whole domain.

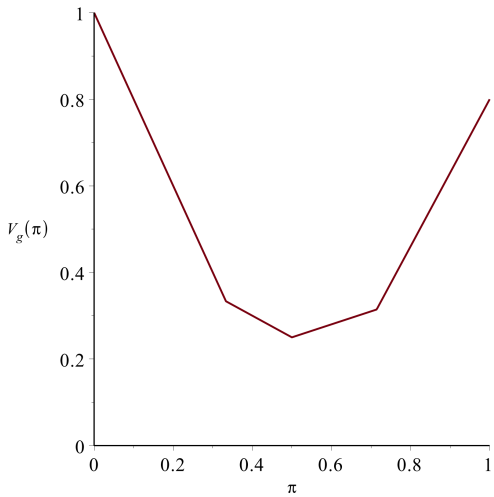

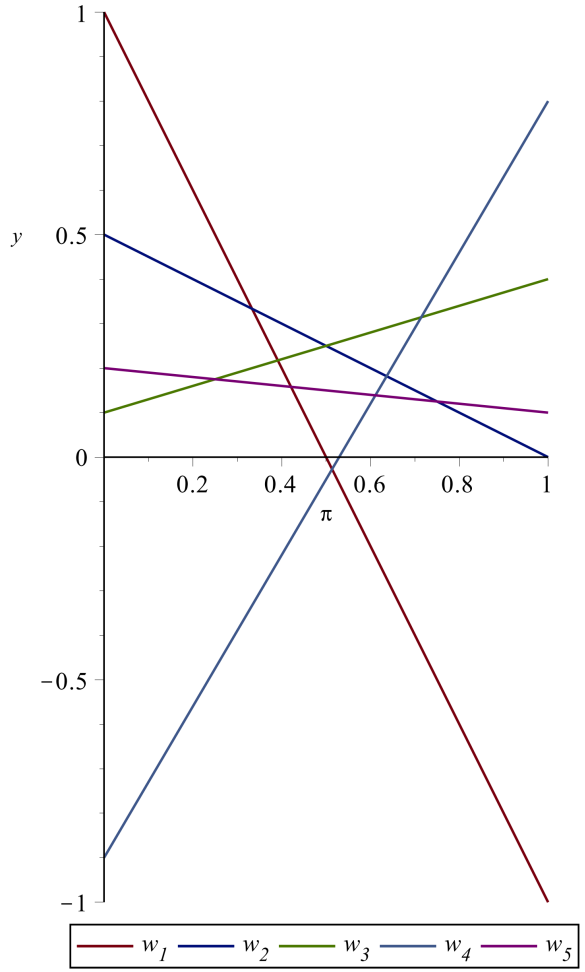

Axiomatization of prior vulnerabilities: Soundness and completeness of V_g wrt. continuous, convex functions

- As usual:
  - A real-valued function f over distributions is **continuous** at π if $f(\pi')$ converges to $f(\pi)$ when π' converges to π.

    (In this case wrt. Def. 11.1 of the Axiom of continuity `CNTY`.)
  - Moreover, f is **continuous everywhere** if it is continuous at all distributions.
- Our first result is that any g-vulnerability V_g for g in $\mathbb{G}\mathcal{X}$ is convex and continuous.

Theorem (11.4)

If g is a gain function in $\mathbb{G}\mathcal{X}$, then the vulnerability V_g is convex and continuous.

Proof. Given in the textbook. □

Axiomatization of prior vulnerabilities: Soundness and completeness of V_g wrt. continuous, convex functions

- Our second result is the converse, that `CVX` and `CNTY` together characterize g-vulnerabilities.

  More precisely, we show that any convex and continuous function over $\mathbb{D}\mathcal{X}$ is expressible as a V_g for some g.

Theorem (11.5)

Let $f : \mathbb{D}\mathcal{X} \to \mathbb{R}_{\geq 0}$ be convex and continuous. Then f is realized as a g-vulnerability V_g for some gain function g in $\mathbb{G}\mathcal{X}$.

Proof. Given in the textbook. □

Axiomatization of prior vulnerabilities: Soundness and completeness of V_g wrt. continuous, convex functions

- The justifications we have so far provided for the axioms of `CVX` and `Q-CVX` might not strike us as very intuitive at first.

  But they can be justified as natural consequences of fundamental axioms relating prior and posterior vulnerabilities, and specific choices for constructing $\widehat{\mathbb{V}}$.

  We'll now address those relationships in detail.

Axiomatization of posterior vulnerabilities

Axiomatization of posterior vulnerabilities

- We will now consider:
  - axioms for posterior vulnerabilities alone,
  - and axioms that relate posterior and prior vulnerabilities to each other.
- We consider three of them, and investigate how different definitions of posterior vulnerabilities shape their interactions.

Axiomatization of posterior vulnerabilities

- The first property that we consider states that a prior π and the point hyper $[\pi]$ are equally good for the adversary.

Definition (11.6 – Axiom of noninterference (NI))

The vulnerability of a point hyper equals the vulnerability of the unique inner of that hyper: for all distributions π in $\mathbb{D}\mathcal{X}$ we have

$$\widehat{\mathbb{V}}[\pi] \quad = \quad \mathbb{V}(\pi) \quad .$$

- We can see why the above is called the axiom of noninterference if we recall Def. 4.14 which states that the abstract channel of a **noninterfering channel** is the mapping

$$\pi \mapsto [\pi] \ .$$

The axiom states that a noninterfering channel should leak nothing when its leakage is measured by comparing $\widehat{\mathbb{V}}$ and $\mathbb{V}$.

Axiomatization of posterior vulnerabilities

- The second axiom we consider is an analog of the famous Data-Processing Inequality (DPI) for mutual information, which:
  - states that if $X \rightarrowtriangle Y \rightarrowtriangle Z$ forms a Markov chain, then $I(X;Y) \geq I(X;Z)$, and
  - can be interpreted as *"post-processing can only destroy information"*.
- We formalize our version of DPI as follows.

Definition (11.7 – Axiom of data-processing inequality (DPI))

Post-processing does not increase vulnerability: for all distributions π in $\mathbb{D}\mathcal{X}$ and (conformal) channel matrices C,R we have

$$\widehat{\mathbb{V}}[\pi \triangleright \mathsf{C}] \quad \geq \quad \widehat{\mathbb{V}}[\pi \triangleright \mathsf{CR}] \quad .$$

(Notice that since the DPI axiom uses cascading, it cannot be expressed in terms of abstract channels: they must be concrete.)

Axiomatization of posterior vulnerabilities

- The `DPI` axiom can be interpreted as follows.

  Consider a secret X:

  1. is fed into a (concrete) channel C, and
  2. then the produced output is post-processed by being fed into another (concrete) channel R.

  Now consider two adversaries A and A' such that:

  - A can observe only the output of channel C, so A's posterior knowledge about the secret is given by the hyper $[\pi \triangleright \mathsf{C}]$.
  - A' can observe only the output of the cascade $\mathsf{C}' = \mathsf{CR}$, so A''s posterior knowledge about the secret is given by $[\pi \triangleright \mathsf{C}']$.

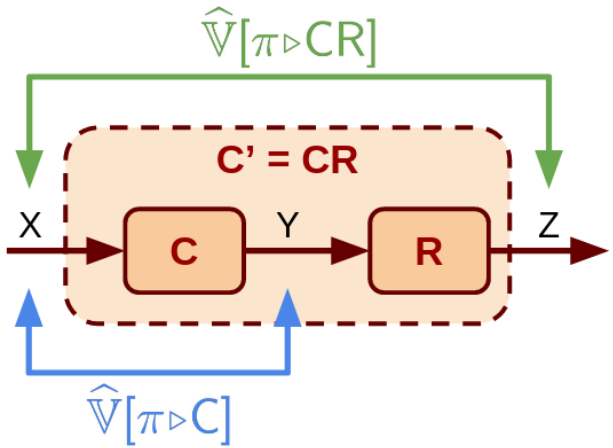

  Now from A's knowledge it is always possible to reconstruct what A' knows, but the converse is not necessarily true.

Axiomatization of posterior vulnerabilities

- To see that, note that:
  - A can use π and C to compute $[\pi \triangleright \mathsf{CR}']$ for any R', including the particular R used by A'.
  - On the other hand, A' knows only π and C' and, in general, the decomposition of C' into a cascade of two channels is not unique.

    More precisely, there may be many pairs C'' and R'' of matrices satisfying $\mathsf{C}' = \mathsf{C}''\mathsf{R}''$.

    Hence it is not always possible for A' to uniquely recover C from C' and then compute $[\pi \triangleright \mathsf{C}]$.

Given that asymmetry, the `DPI` formalizes that a vulnerability $\widehat{\mathbb{V}}$ should not evaluate A's information as any less of a threat than A''s.

Axiomatization of posterior vulnerabilities

- The third property we consider is that by observing the output of a channel an adversary cannot lose information about the secret.

  In the worst case, the output can simply be ignored if it is not useful.

- We formalize that property as follows.

Definition (11.8 – Axiom of monotonicity (MONO))

Pushing a prior through a channel does not decrease vulnerability: for all distributions π in $\mathbb{D}\mathcal{X}$ and channels C we have

$$\widehat{\mathbb{V}}[\pi \triangleright C] \quad \geq \quad \mathbb{V}(\pi) \quad .$$

- The monotonicity axiom has two direct consequences on leakage:
  - additive leakage is always non-negative, and
  - multiplicative leakage is never smaller than 1.

Axiomatization of posterior vulnerabilities

- Notice that in the MONO axiom we are adopting a static perspective on leakage, meaning that
  - even if vulnerability may decrease on some specific output of a channel,
  - that decrease will be compensated for by vulnerability's increase on other potential outputs.

Axiomatization of posterior vulnerabilities: Possible definitions of posterior vulnerabilities

- Having introduced the axioms of
  - `NI`,
  - `DPI`,
  - `MONO`,

  we consider how posterior vulnerabilities can be defined so as to respect them.
- Recall that in the case of prior vulnerabilities, the axioms considered
  - `CVX`,
  - `CNTY`

  uniquely determined the set of prior g-vulnerabilities.
- In contrast, in the case of posterior vulnerability the axioms considered so far do not uniquely determine the set of posterior g-vulnerabilities.

  For that reason, in the following we shall:
  - consider alternative definitions of posterior vulnerabilities, and
  - discuss the interrelations of axioms each definition induces.

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- As we have seen, the posterior versions of:
  - Bayes vulnerability,
  - g-vulnerability,
  - Shannon entropy, and
  - guessing entropy

  are all defined as the expectation of the corresponding prior measures over the (hyper-)distribution of posterior distributions.
- That definition of posterior vulnerability as expectation is now formalized.

Definition (11.9 – Axiom of averaging (AVG))

The vulnerability of a hyper is the expected value wrt. the outer distribution of the vulnerabilities of its inners: for all hyper-distributions Δ in $\mathbb{D}^2\mathcal{X}$ we have

$$\widehat{\mathbb{V}}\Delta \quad = \quad \mathcal{E}_{\Delta}\mathbb{V} \quad ,$$

where the hyper $\Delta{:}\,\mathbb{D}\mathcal{X}$ typically results from $\Delta{=}[\pi\triangleright C]$ for some π, C.

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- Recall our notation that

$$\mathcal{E}_\alpha F$$

  is the expected value of real-valued function F over a distribution α.

- So, in the definition of the `AVG` axiom, when we have

$$\widehat{\mathbb{V}}\Delta \quad = \quad \mathcal{E}_\Delta \mathbb{V}$$

  that means that

  - F is $\mathbb{V}$, a function from inner to real, and
  - α is Δ, a distribution on inners.

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- By imposing AVG on a prior/posterior pair $(\mathbb{V}, \widehat{\mathbb{V}})$ of vulnerabilities, we can uncover a series of interesting relations among other axioms.
- The figure shows equivalences of axioms given AVG.

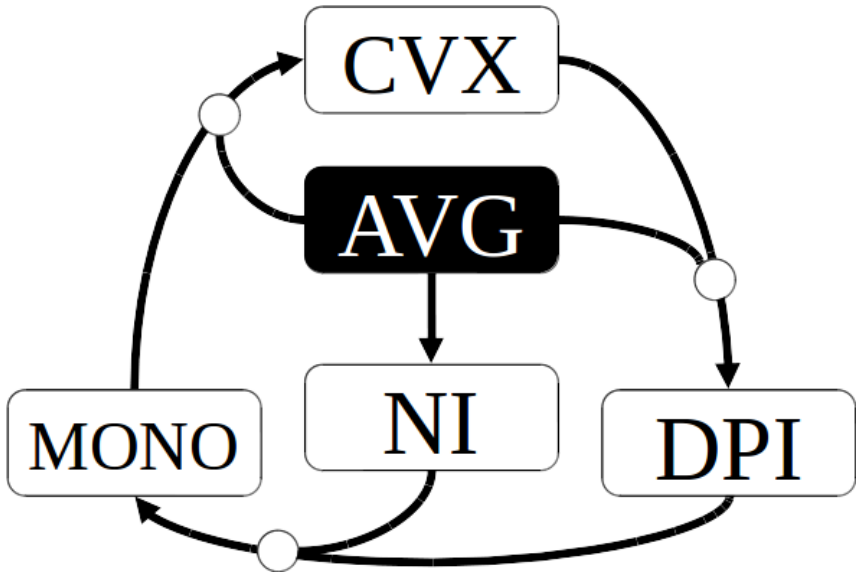

Merging arrows indicate joint implication.

E.g., given AVG we have that CVX implies DPI.

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- We begin by showing that AVG implies NI.

Theorem (11.10 – AVG $\Rightarrow$ NI)

If a pair of prior/posterior vulnerabilities $(\mathbb{V}, \widehat{\mathbb{V}})$ satisfies AVG, then it also satisfies NI.

Proof. If AVG is assumed then for any prior π we have $\widehat{\mathbb{V}}[\pi] = \mathcal{E}_{[\pi]}\mathbb{V} = \mathbb{V}(\pi)$, since $[\pi]$ is a point hyper. □

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- Second, we show that the axioms of NI and DPI, taken together, imply MONO.

Theorem (11.11 – NI ∧ DPI ⇒ MONO)

If a pair of prior/posterior vulnerabilities $(\mathbb{V}, \widehat{\mathbb{V}})$ satisfies NI and DPI, then it also satisfies MONO.

Proof. Take any π, C, and recall that $\mathbb{1}$ denotes the noninterfering channel with only one column and as many rows as the columns of C. Then we reason

$$
\begin{array}{lll}
 & \widehat{\mathbb{V}}[\pi \triangleright \mathsf{C}] & \\
\geq & \widehat{\mathbb{V}}[\pi \triangleright \mathsf{C}\mathbb{1}] & \text{“by DPI”} \\
= & \widehat{\mathbb{V}}[\pi \triangleright \mathbb{1}] & \text{“}\mathsf{C}\mathbb{1} = \mathbb{1}\text{”} \\
= & \widehat{\mathbb{V}}[\pi] & \text{“definition } [- \triangleright -]\text{”} \\
= & \mathbb{V}(\pi) \quad . & \text{“by NI”}
\end{array}
$$

□

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- Third, we show that the axioms AVG and MONO together imply CVX.

Theorem (11.12 – AVG ∧ MONO ⇒ CVX)

If a pair of prior/posterior vulnerabilities $(\mathbb{V}, \widehat{\mathbb{V}})$ satisfies AVG and MONO, then it also satisfies CVX.

Proof. Let $\mathcal{X}=\{x_1, \ldots, x_N\}$ be finite, and let π^1 and π^2 be distributions over $\mathcal{X}$.

Let $0 < a < 1$ so that also $\pi^3 = a\pi^1 + (1-a)\pi^2$ is a distribution on $\mathcal{X}$.

(In case $a=0$ or $a=1$ we would have $\pi^3 = \pi^1$ or $\pi^3 = \pi^2$, respectively, and convexity would follow trivially.)

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- **Proof. (Continued)** Define C^* to be the two-column channel matrix

$$\mathsf{C}^* \quad = \quad \begin{bmatrix} a\pi_1^1/\pi_1^3 & (1-a)\pi_1^2/\pi_1^3 \\ \vdots & \vdots \\ a\pi_n^1/\pi_n^3 & (1-a)\pi_n^2/\pi_n^3 \\ \vdots & \vdots \\ a\pi_N^1/\pi_N^3 & (1-a)\pi_N^2/\pi_N^3 \end{bmatrix}$$

in which there is a row for each element x_n of $\mathcal{X}$ with $\pi_n^3 \neq 0$.

(Note that if $\pi_n^3 = 0$ for some n, then x_n is in the support of neither π^1 nor π^2 (since $0<a<1$), and wlog. we can remove element x_n from both priors and from the channel matrix C^* above.)

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- **Proof. (Continued)** By pushing π^3 through C^* we obtain the hyper $[\pi^3 \triangleright \mathsf{C}^*]$ with outer distribution $(a, 1-a)$, and associated inners π^1 and π^2.

  Since `AVG` is assumed, we have

$$\widehat{\mathbb{V}}[\pi^3 \triangleright \mathsf{C}^*] \quad = \quad a\mathbb{V}(\pi^1) + (1-a)\mathbb{V}(\pi^2) \quad . \tag{1}$$

  But note that by `MONO`, we also have

$$\widehat{\mathbb{V}}[\pi^3 \triangleright \mathsf{C}^*] \quad \geq \quad \mathbb{V}(\pi^3) \quad = \quad \mathbb{V}(a\pi^1 + (1-a)\pi^2) \quad . \tag{2}$$

  Taking (1) and (2) together, we obtain `CVX`. □

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- Finally, we show that the axioms `AVG` and `CVX` together imply `DPI`.
- For that, we will need the following lemma.

Lemma (11.13)

Let $X \twoheadrightarrow Y \twoheadrightarrow Z$ form a Markov chain whose triply joint distribution is given by $p(x,y,z) = p(x)\,p(y|x)\,p(z|y)$ for all (x,y,z) in $\mathcal{X}\times\mathcal{Y}\times\mathcal{Z}$.

Then we have that $\sum_y p(y|z)\,p(x|y) = p(x|z)$ for all x, y, z.

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- **Proof.** First we note that the probability of z depends only on the probability of y, and not x, and so $p(z|x,y) = p(z|y)$ for all x, y, z.

  Then we can use the fact that

$$p(y,z)\,p(x,y) \quad = \quad p(x,y,z)\,p(y) \tag{3}$$

  to reason

$$\begin{aligned}
& \textstyle\sum_y \; p(y|z)p(x|y) & \\
= \quad & \textstyle\sum_y \; {}^{p(y,z)}\!/_{p(z))} \times {}^{p(x,y)}\!/_{p(y)} & \text{“by definition of conditional”} \\
= \quad & \textstyle\sum_y \; p(x,y,z)p(y) \,/\, p(z)p(y) & \text{“by (3)”} \\
= \quad & \textstyle\sum_y \; p(x,y \mid z) & \text{“by definition of conditional”} \\
= \quad & p(x|z) \quad . & \text{“by marginalization”}
\end{aligned}$$

  □

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- Now we present another relation among axioms.

Theorem (11.14 – AVG $\wedge$ CVX $\Rightarrow$ DPI)

If a pair of prior/posterior vulnerabilities $(\mathbb{V}, \widehat{\mathbb{V}})$ satisfies AVG and CVX, then it also satisfies DPI.

Proof. Let $\mathcal{X}$, $\mathcal{Y}$ and $\mathcal{Z}$ be sets of values. Let π be a prior on $\mathcal{X}$, and let C be a (concrete) channel from $\mathcal{X}$ to $\mathcal{Y}$, and R be a (concrete) channel from $\mathcal{Y}$ to $\mathcal{Z}$. Note that the cascade CR of channels C and R is a channel from $\mathcal{X}$ to $\mathcal{Z}$.

Let $p(x, y, z)$ be the triply joint distribution defined $p(x, y, z) = \pi_x \mathsf{C}_{x,y} \mathsf{R}_{y,z}$ for all (x, y, z) in $\mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$. By construction, this distribution has the property that the probability of z depends only on the probability of y, and not x, that is that $p(z \,|\, x, y) = p(z \,|\, y)$.

Axiomatization of posterior vulnerabilities: Posterior vulnerability as expectation

- **Proof. (Continued)** Note that, by pushing prior π through channel C, we obtain hyper $[\pi \triangleright \mathsf{C}]$, in which the outer distribution on y is $p(y)$, and the inners are $p_{X|y}$.

  Thus we can reason

$$
\begin{array}{lll}
 & \widehat{\mathbb{V}}[\pi \triangleright \mathsf{C}] & \\
= & \sum_y \; p(y)\mathbb{V}(p_{X|y}) & \text{“by } \texttt{AVG}\text{”} \\
= & \sum_y \; \left(\sum_z p(z)p(y|z)\right)\mathbb{V}(p_{X|y}) & \text{“by marginalization”} \\
= & \sum_z \; p(z)\sum_y p(y|z)\mathbb{V}(p_{X|y}) & \text{“moving constants wrt. the sum”} \\
\geq & \sum_z \; p(z)\,\mathbb{V}\Big(\sum_y \; p(y|z)p_{X|y}\Big) & \text{“by } \texttt{CVX}\text{”} \\
= & \sum_z \; p(z)\,\mathbb{V}(p_{X|z}) & \text{“by Lem. 11.13”} \\
= & \widehat{\mathbb{V}}[\pi \triangleright \mathsf{CR}] \quad . & \text{“by } \texttt{AVG}\text{”}
\end{array}
$$

□

Axiomatization of posterior vulnerabilities: Posterior vulnerability as maximum

- An important consequence of AVG is that an observable's happening with very small probability will have a negligible effect on $\widehat{\mathbb{V}}$, even if that observable reveals the secret completely.
- If that is not acceptable, an alternative approach is to consider the maximum (rather than average) information that can be obtained from any single output of the channel –produced with nonzero probability– no matter how small that probability might be.
- This conservative approach represents a defender who worries about the worst possible amount of threat to the secret.

Axiomatization of posterior vulnerabilities: Posterior vulnerability as maximum

- The definition of posterior vulnerability in those terms is now formalized.

Definition (11.15 – Axiom of maximum (MAX))

The vulnerability of a hyper is the maximum value among the vulnerabilities of the inners in the support of its outer: thus for all hypers Δ we have

$$\widehat{\mathbb{V}}\Delta \quad = \quad \max_{\lceil \Delta \rceil} \mathbb{V} \quad ,$$

where the hyper $\Delta : \mathbb{D}\mathcal{X}$ typically results from $\Delta = [\pi \triangleright C]$ for some π, C.

Note that in the definition above:

- We are writing $\max_S f$ for the maximum value of function f over arguments taken from set S. Here the function is $\mathbb{V}$ on distributions, and the arguments are the inners of Δ, equivalently the elements of the set $\lceil \Delta \rceil$.
- We take the support of the outer distribution because we ignore the vulnerability of inners that cannot happen (i.e. that have probability zero).

Axiomatization of posterior vulnerabilities: Posterior vulnerability as maximum

- By imposing MAX on a prior/posterior pair $(\mathbb{V}, \widehat{\mathbb{V}})$ of vulnerabilities, we can derive relations among other axioms, just as we did for AVG.
- But they are different relations.
  - The symmetry among CVX, MONO and DPI is broken under MAX.
  - Although the axioms of MONO and DPI are still equivalent (a result that we shall soon demonstrate), they are weaker than the axiom of CVX.
- The symmetry can be recovered, however, if we substitute AVG with MAX.
- The figure shows equivalences of axioms given MAX (proven in the textbook).

  Merging arrows indicate joint implication.

  E.g., given MAX we have that Q-CVX implies DPI.

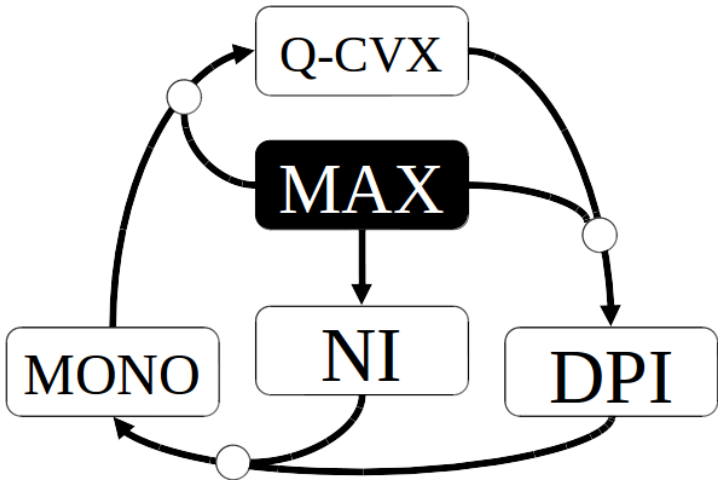

Appendix: Applications of axiomatization to understanding leakage measures

Applications of axiomatization to understanding leakage measures

- The relationships we have uncovered among axioms helps us better to understand the multitude of possible leakage measures one can adopt.

  E.g., what V_g to use, and what version of leakage –additive or multiplicative– to employ

- We now give two examples of how the axiomatization of g-vulnerabilities and leakage can be useful in different scenarios.

- A first instance concerns the robustness of the refinement relation ($\sqsubseteq$).

Applications of axiomatization to understanding leakage measures

- Example 2 Recall that refinement is sound (Thm. 9.11) with respect to the strong g-leakage ordering ($\geq_{\mathbb{G}}$).
  - Still, we might worry that refinement implies a leakage ordering only with respect to g-leakage.
  - That would leave open the possibility that the leakage ordering might conceivably fail for some yet-to-be-defined leakage measure.

  But notice that we have shown

  - $\texttt{AVG} \wedge \texttt{MONO} \Rightarrow \texttt{CVX}$ (Thm. 11.12),
  - $\texttt{AVG} \wedge \texttt{CVX} \Rightarrow \texttt{DPI}$ (Thm. 11.14),

  which, together, show that if the hypothetical new leakage measure

  - is defined using AVG, and
  - never gives negative leakage,

  then it also satisfies the data-processing inequality DPI.

  Hence refinement would be sound for the new leakage measure as well. ◁

Applications of axiomatization to understanding leakage measures

- Another application concerns the possibility of negative leakage for some information measures.

- Example 3 Consider Rényi entropy which, is a set of entropy measures that has been used in the context of quantitative information flow.

  The set is defined by

$$H_\alpha(\pi) \quad = \quad \frac{1}{1-\alpha} \log_2 \left(\sum_{x \in \mathcal{X}} \pi_x^\alpha \right)$$

  for $0 \leq \alpha \leq \infty$ (taking limits in the cases of α=1, which gives Shannon entropy, and $\alpha = \infty$, which gives min-entropy).

Applications of axiomatization to understanding leakage measures

- Example 3 (Continued)

  It would be natural to use Rényi entropy to introduce a set of leakage measures by defining posterior Rényi entropy $\widehat{H}_\alpha$ (notated by analogy with $\widehat{\mathbb{V}}$) using AVG and defining Rényi leakage by

  $$\mathcal{L}_\alpha(\pi, C) \quad = \quad H_\alpha(\pi) - \widehat{H}_\alpha[\pi \triangleright C] \quad .$$

  However, it turns out that H_α is not concave for $\alpha > 2$.

  Therefore, by the dual version of Thm. 11.12 (AVG+MONO $\Rightarrow$ CVX), we find that Rényi leakage $\mathcal{L}_\alpha$ for $\alpha > 2$ would sometimes be negative.

Applications of axiomatization to understanding leakage measures

- Example 3 (Continued)

The figure shows how the nonconcavity of min-entropy H_∞ can cause posterior min-entropy to be greater than prior min-entropy, giving negative min-entropy leakage.

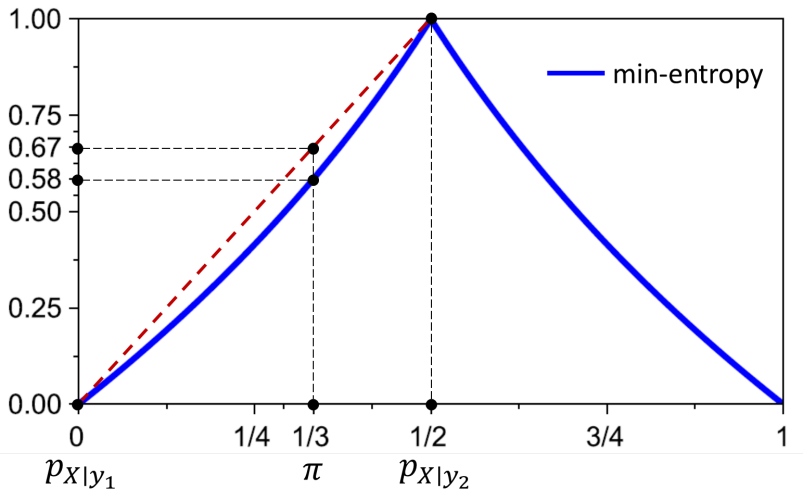

Pushing prior $\pi := (1/3, 2/3)$ through concrete channel C defined

$$\mathsf{C} := \begin{pmatrix} 0 & 1 \\ 1/2 & 1/2 \end{pmatrix}$$

gives hyper $[\pi \triangleright \mathsf{C}]$ with outer $(1/3, 2/3)$ and the two inners $p_{X|y_1} = (0, 1)$ and $p_{X|y_2} = (1/2, 1/2)$.

Thus $H_\infty(\pi) = -\log_2 2/3 \approx 0.58$ and $\widehat{H}_\infty[\pi \triangleright \mathsf{C}] = 1/3 \cdot 0 + 2/3 \cdot 1 \approx 0.67$.

Applications of axiomatization to understanding leakage measures

- Example 3 (Continued)

  This problem is avoided, however, if we do not define posterior min-entropy via `AVG`, but instead by using

$$\widehat{H}_\infty[\pi \triangleright C] := -\log_2 V_1[\pi \triangleright C] \ .$$

◁