The Dalenius perspective

Mário S. Alvim
(msalvim@dcc.ufmg.br)

Quantitative Information Flow

DCC-UFMG
(2021/1)

Introduction

- Now we'll study the effect on information leakage of correlations among different secrets.
- Given
  - secret X,
  - prior π, and
  - channel C,

  our concern has been the leakage of X caused by C.

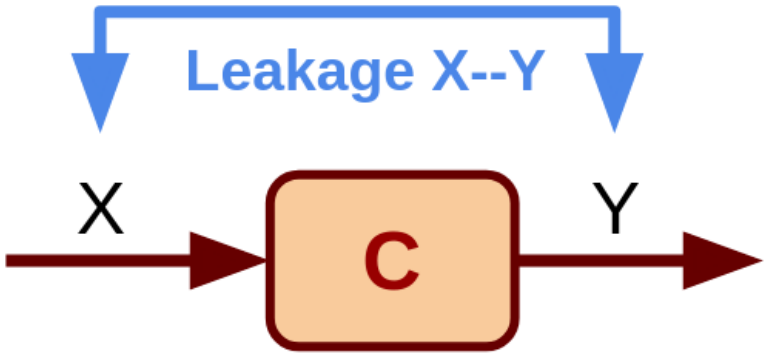

Introduction

- But what if there is another secret Z, which
  1. apparently has nothing to do with C, but
  2. is correlated with X via some joint distribution

$$J : \mathbb{D}(\mathcal{Z} \times \mathcal{X})$$

known to the adversary?

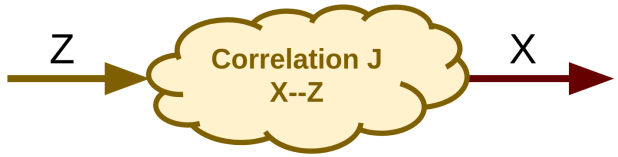

- Then C will also leak information about Z!

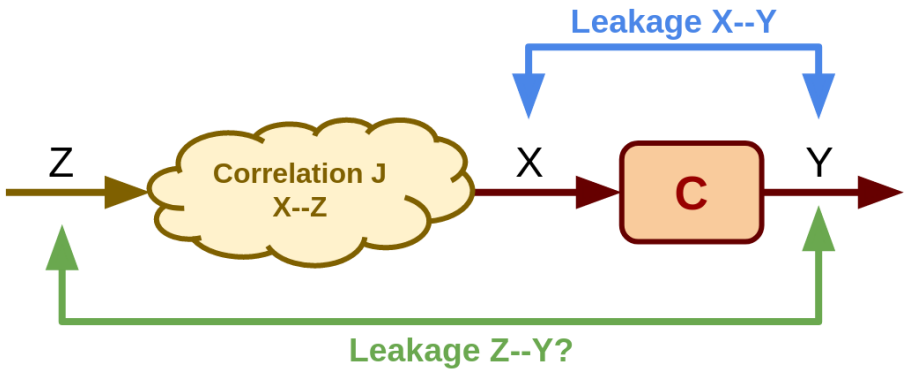

Even if the designer isn't aware of:

- the existence of Z, or
- its correlation via J with X!

- How can we defend against Z's being leaked in this "accidental" way?

Introduction

- Example 1 Suppose that a number of beneficiaries
  1. reside in a region comprising three counties A, B, and C, and
  2. have a variety of ages.

Assume a beneficiary is selected at random, and we consider as secrets both

- the beneficiary's county Z, and
- the beneficiary's age X.

Suppose further that a census has been taken and the following table of "macrostatistics" (from Dalenius) has been published:

Number of beneficiaries by county and age

County	Age class			
	Under 65	65–69	70–74	75 & over
A	3	15	11	8
B	7	60	34	20
C	0	4	0	0

Introduction

- Example 1 (Continued)

We can normalize the table by dividing each entry by the total number of beneficiaries (162), obtaining a joint distribution.

We express such a joint as a matrix J indexed by $\mathcal{Z}$ (county) and $\mathcal{X}$ (age).

J	Under 65	65–69	70–74	75 & over
A	$^{3}/_{162}$	$^{15}/_{162}$	$^{11}/_{162}$	$^{8}/_{162}$
B	$^{7}/_{162}$	$^{60}/_{162}$	$^{34}/_{162}$	$^{20}/_{162}$
C	0	$^{4}/_{162}$	0	0

By summing the rows and the columns of J we can compute

the marginal on Z: $\rho = (^{37}/_{162},\ ^{121}/_{162},\ ^{4}/_{162})$, and

the marginal on X: $\pi = (^{10}/_{162},\ ^{79}/_{162},\ ^{45}/_{162},\ ^{28}/_{162})$.

- Example 1 (Continued)

Those marginals ρ on Z and π on X themselves are of course already useful to an adversary interested in secrets Z or X separately.

But more worrying is that the correlation means that a channel C leaking information about age will also leak information about county!

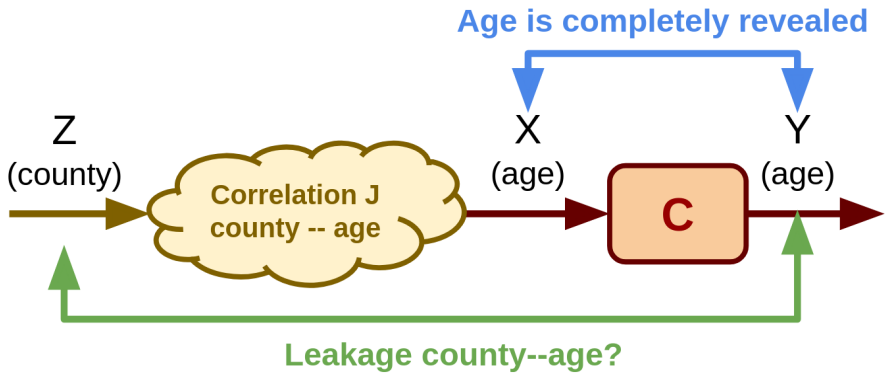

For instance,

- if C reveals that the age is not in the range 65–69, then
- the adversary can deduce that the county isn't C!

J	Under 65	65–69	70–74	75 & over
A	$3/162$	$15/162$	$11/162$	$8/162$
B	$7/162$	$60/162$	$34/162$	$20/162$
C	0	$4/162$	0	0

Introduction

- Note especially that knowledge of the channel C is not necessarily related in any way to knowing there is a correlation J.
  - It is the adversary who knows J.
  - The data partly revealed by C could be via some other study entirely!
- This further complicates the design of a "non-leaky" channel C from X to Y.

  How can we control the leakage our channel C causes about another secret Z (correlated to X), if we might not even know Z even exists?

- We call such "collateral leakage" the **Dalenius leakage** of Z caused by channel C under correlation J.

  Our goal here is to develop the theory of such leakage.

Dalenius scenarios

Dalenius scenarios

- In the scenario of interest here, there is:
  - a channel C that (potentially) leaks information about a secret X, and
  - an adversary who
    - actually cares about a different secret Z, and
    - who knows not only the prior π of X but also the correlation between Z and X.

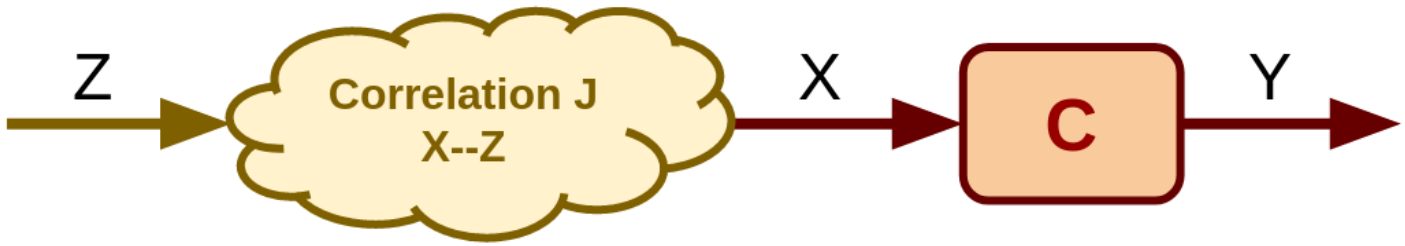

- Given what the adversary might find out about X, she could then go further:
  1. First, her observation of C acting on prior π induces a posterior distribution (an inner) $p_{X|y}$ on X, which is what she learns about X.
  2. Then, with the correlation between Z and X, she can use that $p_{X|y}$ to induce a posterior distribution on Z.

Dalenius scenarios

- We describe the Z-to-X correlation as a joint distribution $J\colon \mathbb{D}(\mathcal{Z}\times\mathcal{X})$.

  We note that its X-marginal is the prior π for C.

- From J, we get marginal distributions
  - $\pi\colon \mathbb{D}\mathcal{X}$, representing the adversary's prior knowledge about X, and
  - $\rho\colon \mathbb{D}\mathcal{Z}$, representing her prior knowledge about Z.

- Moreover, when we express J as a matrix J, we can factor it via

$$\mathsf{J} = \rho \triangleright \mathsf{B} \ ,$$

  for some (stochastic) matrix $\mathsf{B}\colon \mathcal{Z} \rightarrowtriangle \mathcal{X}$.

  More precisely:
  - $\rho_z \mathsf{B}_{z,x} = \mathsf{J}_{z,x}$ for all z and x or, and
  - each row $\mathsf{B}_{z,-}$ of B is found simply by normalizing row $\mathsf{J}_{z,-}$.

Dalenius scenarios

- Now suppose we take a concrete realization C of our abstract channel C.

  Recall that by the theory of cascading, we know that

$$
\begin{aligned}
p(y|z) &= \sum_x p(x, y \mid z) && \text{(marginalization)} \\
&= \sum_x p(x \mid z) p(y \mid x, z) && \text{(chain rule)} \\
&= \sum_x p(x \mid z) p(y \mid x) && (y \text{ is indep. of } z \text{ given } x) \\
&= \sum_x \mathsf{B}_{z,x} \mathsf{C}_{x,y} && \text{(channel notation)} \\
&= (\mathsf{BC})_{z,y} && \text{(def. of matrix multiplication)}
\end{aligned}
$$

  That leads us to the following remarkable conclusion:

  The adversary's "Dalenius knowledge" about Z that results from running C and revealing Y is simply the hyper-distribution $[\rho \triangleright \mathsf{BC}]$!

Dalenius scenarios

- That justifies the following definition.

Definition (10.1 – Dalenius g-vulnerability)

Suppose we are given a channel C on X and a joint distribution $J: \mathbb{D}(\mathcal{Z}\times\mathcal{X})$ that represents an adversary's knowledge of a correlation between $\mathcal{Z}$ and $\mathcal{X}$. Let J be a matrix realization of J, and factor it into marginal distribution $\rho: \mathbb{D}\mathcal{Z}$ and stochastic matrix B: $\mathcal{Z} \to\!\!\triangleright \mathcal{X}$, i.e. so that $\mathsf{J} = \rho \triangleright \mathsf{B}$. Let C be a concrete realization of C.

Then, for any gain function $g: \mathbb{G}\mathcal{Z}$, the **Dalenius g-vulnerability** of J and C is defined

$$V_g^D(J, C) \;:=\; V_g[\,\rho \triangleright \mathsf{BC}\,] \quad .$$

The special case when g is g_{id} is called **Dalenius Bayes vulnerability**, written V_1^D.

- Note that the definition above does not depend on the precise realizations of J and C, given Def. 9.6 of abstract refinement.

Dalenius scenarios

- The definition implies the following important result, which shows that refinement ($\sqsubseteq$) is also equivalent to the **strong Dalenius vulnerability ordering**.

  More precisely, the result shows that

$$C \sqsubseteq D$$

iff

$$V_g^D(J, C) \quad \geq \quad V_g^D(J, D) \quad \text{for all correlations } J \text{ and gain functions } g.$$

Theorem (10.2)

For any (abstract) channels C and D on X, we have $C \sqsubseteq D$ iff the Dalenius g-vulnerability of D never exceeds that of C, for any correlation J between X and some Z and any gain function g.

Dalenius scenarios

- **Proof.** Write $(\sqsubseteq^D)$ for the strong Dalenius g-vulnerability order (i.e. that for all g and J we have $V_g^D(J, C) \geq V_g^D(J, D)$).

  Recall that

  $$(\sqsubseteq_\circ) \quad = \quad (\sqsubseteq) \quad = \quad (\sqsubseteq_{\mathbb{G}}) \,.$$

  Then $(\sqsubseteq^D)$ implies $(\sqsubseteq_{\mathbb{G}})$ because $(\sqsubseteq_{\mathbb{G}})$ is the special case of $(\sqsubseteq^D)$ where $Z{=}X$ and J as a matrix is $\ulcorner \pi \urcorner$. (This is where the Coriaceous theorem 9.12 is used.)

  For the opposite direction, that $(\sqsubseteq_\circ)$ implies $(\sqsubseteq^D)$, let concrete channels C and D realize C and D, respectively. Note that although C and D must have the same input type $\mathcal{X}$, they can have different output types say $\mathcal{Y}^C$ and $\mathcal{Y}^D$.

  Now from Def. 9.5 (concrete refinement) we have that $\mathsf{C} \sqsubseteq_\circ \mathsf{D}$ gives $\mathsf{CR}{=}\mathsf{D}$ for some R of type $\mathcal{Y}^C \to \mathcal{Y}^D$, whence immediately $(\mathsf{BC})\mathsf{R} = \mathsf{BD}$ and thus $\mathsf{BC} \sqsubseteq_\circ \mathsf{BD}$, giving our desired inequality $V_g[\,\rho \triangleright \mathsf{BC}\,] \geq V_g[\,\rho \triangleright \mathsf{BD}\,]$ for all ρ, g — which is $(\sqsubseteq^D)$ by Def. 10.1 (Dalenius g-vulnerability). □

Dalenius scenarios

- Example 2 Suppose that we have a 2-bit secret X and a channel C, expressed concretely as C of type $\mathcal{X} \to \{0,1\}$, that leaks the least-significant bit of X:

C	0	1
00	1	0
01	0	1
10	1	0
11	0	1

.

Suppose further that there is a 1-bit secret Z that is correlated with X according to the joint-distribution matrix J^{ZX}:

J^{ZX}	00	01	10	11
0	$1/8$	$1/16$	$1/4$	$1/16$
1	$1/16$	$1/4$	$1/16$	$1/8$

.

As defenders of X we might be completely unaware of that data Z and –as such– have no interest in its security, and so take no measures to ensure it.

What leakage do J^{ZX} and C imply?

Dalenius scenarios

- Example 3 From J^{ZX} we obtain marginal distributions ρ, π on $\mathcal{Z}, \mathcal{X}$:

	0	1
ρ	$1/2$	$1/2$

and

	00	01	10	11
π	$3/16$	$5/16$	$5/16$	$3/16$

.

Now we find the posterior Bayes vulnerability of X due to C's acting on π.

By Thm. 5.15 we the posterior Bayes vulnerability as the sum of the column maximums of J^{XY}:

J^{XY}	0	1
00	$3/16$	0
01	0	$5/16$
10	$5/16$	0
11	0	$3/16$

$$V_1[\pi \triangleright C] \quad = \quad 5/16 + 5/16 \quad = \quad 5/8 \quad .$$

Since X's prior Bayes vulnerability is

$$V_1(\pi) = 5/16 \ ,$$

we observe that X is leaked by C as what we expected.

Dalenius scenarios

- Example 3 (Continued)

But how is Z, which indeed we might not even have heard of, affect by C?

Following Def. 10.1 (Dalenius g-vulnerability), we first get B by normalizing the rows of J^{ZX} and then we compute the cascade BC:

B	00	01	10	11
0	$1/4$	$1/8$	$1/2$	$1/8$
1	$1/8$	$1/2$	$1/8$	$1/4$

$\cdot$

C	0	1
00	1	0
01	0	1
10	1	0
11	0	1

$=$

BC	0	1
0	$3/4$	$1/4$
1	$1/4$	$3/4$

,

from which the Dalenius Bayes vulnerability of Z is

$$V_1^D(J, C) \quad = \quad V_1[\,\rho \triangleright \mathrm{BC}\,] \quad = \quad 3/4 \quad ,$$

since ρ is $(1/2, 1/2)$.

Dalenius scenarios

- Example 3 (Continued)

Since Z's prior Bayes vulnerability is $1/2$, we see that C together with $J : \mathbb{D}(\mathcal{Z} \times \mathcal{X})$ causes leakage of Z — which is what one might not expect.

But notice that:

- defenders of X don't know that (because they are ignorant of Z); and
- defenders of Z don't know that (because they are ignorant of C).

The whole Dalenius issue is, then:

"How can we defend secrets like Z from attacks on X?"

We'll return to that question soon.

But our very next topic is the fact that the existence of Dalenius threats can be used to justify our definition of refinement.

(For that we return to the idea of compositional closure.) ◁

Compositional closure for Dalenius contexts

Compositional closure for Dalenius contexts

- We now come back to why refinement ($\sqsubseteq$) has to be so complicated.

  We'll give a technical argument for the "inevitability" of our definition of ($\sqsubseteq$).

- Example 4 Imagine that
  - a customer has asked for channel S (the specification), and
  - the developer has delivered channel K (the putative implementation).

  Suppose further that

$$V_g[\pi \triangleright S] \quad \geq \quad V_g[\pi \triangleright K]$$

  - for all π's and
  - almost all g's.

  The developer insists:

> *"Rest assured that I have covered enough g's to describe any adversary one might meet in practice."*

  (E.g., maybe he considered all n-guess vulnerabilities $V_n(-)$ for any $n \geq 1$.)

Compositional closure for Dalenius contexts

- Example 4 (Continued)

To continue the story, the customer says:

"That's all very well...

But in fact for this prior $\hat{\pi}$ and this gain function $\hat{g}$, encountered where I deployed your K, it turns out that $V_{\hat{g}}[\hat{\pi} \triangleright S] < V_{\hat{g}}[\hat{\pi} \triangleright K]$.

Hence I am not going to pay you for K."

The developer replies by asserting:

"But that $\hat{g}$ you found is ridiculous, a 'monster'.
It's simply unreasonable to worry about it!"

Now, who is right?

- Well, it turns out that the customer is (always) right...

  No matter how "monstrous" his $\hat{\pi}$ and $\hat{g}$ might be!

Compositional closure for Dalenius contexts

- Example 4 (Continued)

To see that, we look at the above scenario slightly more abstractly.

Customers:

- generally will not know the intricacies of cyber-security and *QIF* but
- still they will in some cases, simple ones, know without any doubt that security really has been compromised.

As a good example of that, consider the **strong Bayes vulnerability order**, defined as

$$S \preccurlyeq K \qquad \text{iff} \qquad \text{for every prior } \pi, \quad V_1[\pi \triangleright S] \geq V_1[\pi \triangleright K] \ .$$

If $S \not\preccurlyeq K$ then even the developer would agree that there's a problem.

- That failed inequality means that, for some priors, the probability of guessing the secret exactly, with K, can be strictly more than with S.

Compositional closure for Dalenius contexts

- Example 4 (Continued)

Failing ($\precnsim$) is a blatant failure; the difficulty comes for the customer when:

- $S \precnsim K$ does hold, and
- his problem is revealed only in a more exotic setting, described by a gain function $\hat{g}$ that the developer claims "in practice can never happen".

In this case the customer can no longer appeal to the "any fool can see" argument that a glaring ($\not\approx$) would provide.

But here is how the customer proves that the "rogue" $\hat{g}$ is worrisome.

1. He exhibits a context $\mathcal{F}$ (for "fail") such that

$$\mathcal{F}S \not\approx \mathcal{F}K\,,$$

2. and his argument becomes

"If I run your K in context $\mathcal{F}$, then its failure is revealed in just one guess. And any fool can see that, at least."

Compositional closure for Dalenius contexts

- The technique of the example above is what we call "compositional closure", and it is how we move:
  - from the simple refinement definition ($\preccurlyeq$) that everyone accepts
  - to the (perhaps) more complicated one that everyone needs.

  More precisely,

  1. We start with a definition of ($\preccurlyeq$) that is generally agreed to be reasonable.

     (That of course remains a subjective issue.)

  2. Then the developer consults *QIF* experts (for which he need not involve the customer) and asks

     "What refinement relation ($\sqsubseteq_?$) do I need to use in house so that if I establish $S \sqsubseteq_? I$ then I can be sure that in the field, i.e. for any context $\mathcal{C}$, the customer will always find that $\mathcal{C}S \preccurlyeq \mathcal{C}I$?"

     The developer is in fact asking the *QIF* expert for the compositional closure of ($\preccurlyeq$) wrt. allowed contexts $\mathcal{C}$ from some agreed-upon set $\mathbb{C}$ of contexts.

Compositional closure for Dalenius contexts

- Motivated by the example above, suppose then that we are considering some customer-plausible relation ($\preceq_?$) of refinement.

  The *QIF* expert will use the following procedure to recommend a suitable ($\sqsubseteq_?$) for the developer.

Definition (Compositional closure)

The **compositional closure** of a primitive refinement relation ($\preceq_?$) with respect to set of contexts $\mathbb{C}$ is the unique relation ($\sqsubseteq_?$) such that

($\sqsubseteq_?$) is *safe* for ($\preceq_?$) — for all channels A, B and contexts $\mathcal{C}$ in $\mathbb{C}$ we have that $A \sqsubseteq_? B$ implies $\mathcal{C}A \preceq_? \mathcal{C}B$. (The customer will never be disappointed, no matter how B is deployed.)

($\sqsubseteq_?$) is *necessary* for ($\preceq_?$) — for all channels A, B we have that $A \not\sqsubseteq_? B$ implies there is some context $\mathcal{C}$ in $\mathbb{C}$ such that $\mathcal{C}A \not\preceq_? \mathcal{C}B$. (The developer is not restricting his refinement techniques unnecessarily.)

Compositional closure for Dalenius contexts

- And now we'll pursue a technical justification for the refinement order ($\sqsubseteq$) that we use.

  We'll argue that it is simply that for "Dalenius contexts", the appropriate approach is compositional closure.

- Since the failure of ($\preccurlyeq$) is not tolerated by anyone, we could argue that our ($\sqsubseteq$) is the inevitable –and unique– choice for refinement of channels.

  It accounts for both the programs into which the customer might put them and the Dalenius contexts they might encounter.

Compositional closure for Dalenius contexts: Safety and necessity wrt. Dalenius contexts

- We have just seen that an argument that our definition of refinement ($\sqsubseteq$) could in a sense be made "inevitable" once the set of allowable contexts had been determined.

  Such sets of contexts could include:

  - Sequential programs in which (side-) channels might occur.
  - Those defined Dalenius leakage, for which we can now be precise.

- In fact, for Dalenius contexts, we have proved safety already, i.e., that

$$\text{if} \quad C \sqsubseteq D \qquad \text{then} \qquad V_1^D(J, C) \;\geq\; V_1^D(J, D), \quad \text{for any } J.$$

  That's because Thm. 10.2 proved the stronger result (in its forward direction) for any vulnerability V_g, not just Bayes vulnerability.

Compositional closure for Dalenius contexts: Safety and necessity wrt. Dalenius contexts

- That leaves necessity, for which we have the following theorem.

Theorem (10.4 – Dalenius necessity)

Let C, D be channels such that $C \not\sqsubseteq D$.

Then there is a Dalenius correlation J such that $V_1^D(J, C) \not\succcurlyeq V_1^D(J, D)$.

That is, distinctions made with our refinement order ($\sqsubseteq$) can be reduced to distinctions made with the strong Bayes-vulnerability order ($\preccurlyeq$) via a suitable Dalenius context J.

Compositional closure for Dalenius contexts: Safety and necessity wrt. Dalenius contexts

- **Proof.** Let π and g be the prior and gain function that establish $C \not\sqsubseteq D$, so

$$V_g[\pi \triangleright C] \quad \not\geq \quad V_g[\pi \triangleright D] \quad , \tag{1}$$

and let C: $\mathcal{X} \rightarrowtriangle \mathcal{Y}^C$ and D: $\mathcal{X} \rightarrowtriangle \mathcal{Y}^D$ be concrete channels corresponding to C, D. (Recall from Def. 9.6 (abstract refinement) that whether or not refinement holds is unaffected by which specific corresponding concrete realizations we take.)

With the trace-based formulation of posterior g-vulnerability given in Thm. 5.24, from (1) we have

$$V_g[\pi \triangleright \mathsf{C}] \quad = \quad \max_{\mathsf{S}} \ \mathbf{tr}(\mathsf{G}\, \lceil \pi \rfloor\, \mathsf{CS}) \quad ,$$

where S ranges over strategies from $\mathcal{Y}^C$ to $\mathcal{W}$, and G is the matrix representation of g, and as usual $\lceil \pi \rfloor$ denotes the $\mathcal{X} \times \mathcal{X}$ matrix with π on its diagonal and zeroes elsewhere.

Compositional closure for Dalenius contexts: Safety and necessity wrt. Dalenius contexts

- **Proof. (Continued)** We choose 1-summing joint-distribution matrix J and positive scalar c so that $c\,\mathsf{J} = \mathsf{G}\lceil\pi\rfloor$. (Note that we are assured that we can do this because the proof of Thm. 9.12 (Coriaceous) always constructs a gain function g that is non-negative, with finitely many actions.)

  Continuing, we reason

$$
\begin{array}{lll}
 & V_g[\pi \triangleright \mathsf{C}] & \\
= & \max_{\mathsf{S}}\ \mathbf{tr}(\mathsf{G}\lceil\pi\rfloor\mathsf{CS}) & \text{“from above”} \\
= & \max_{\mathsf{S}}\ \mathbf{tr}(c \times \mathsf{JCS}) & \text{“choice of } c \text{ and J”} \\
= & c \times \max_{\mathsf{S}}\ \mathbf{tr}(\mathsf{I}\lceil\rho\rfloor\mathsf{BCS}) & \text{“factorize J}{=}\rho \triangleright \mathsf{B}\text{; take } c \text{ outside max; identity matrix I”} \\
= & c \times V_1[\rho \triangleright \mathsf{BC}] & \text{“Thm. 5.24”} \\
= & c \times V_1^D(\mathsf{J}, \mathsf{C}) \quad ; & \text{“Def. 10.1”}
\end{array}
$$

  and similarly for D we can show $V_g[\pi \triangleright \mathsf{D}] = c \times V_1^D(\mathsf{J}, \mathsf{D})$.

  Our result now follows from (1) above. □

Compositional closure for Dalenius contexts: Safety and necessity wrt. Dalenius contexts

- Together with safety, Thm. 10.4 (Dalenius necessity) gives us that our refinement ($\sqsubseteq$) is the (unique) compositional closure of ($\preccurlyeq$), that is the strong Bayes-vulnerability order with respect to Dalenius contexts.

Our refinement order is a compositional closure of Bayes vulnerability

Our refinement order ($\sqsubseteq$) is the compositional closure, under the set of Dalenius contexts $\mathbb{C}$, of the strong Bayes-vulnerability order ($\preccurlyeq$).

We assume always that $\mathbb{C}$ includes the identity context and that it is closed under composition, i.e. that if $\mathcal{C}, \mathcal{C}'$ are in $\mathbb{C}$ then so is $\mathcal{C} \circ \mathcal{C}'$.

- This shows that our definition of refinement ($\sqsubseteq$) can be justified by reducing it to the more intuitive strong Bayes-vulnerability order ($\preccurlyeq$).

We now give an example of how this happens in Dalenius contexts.

Compositional closure for Dalenius contexts: Justifying refinement - An example

- **Example 5** Consider the following concrete channels:

C	y_1	y_2	y_3
x_1	$1/2$	$1/2$	0
x_2	$1/2$	0	$1/2$
x_3	0	$1/2$	$1/2$

and

D	z_1	z_2
x_1	$2/3$	$1/3$
x_2	$2/3$	$1/3$
x_3	$1/4$	$3/4$

.

It can be shown that (in Example 9.8 in our textbook) that:

1. D does not refine C, i.e.,
$$\mathsf{C} \not\sqsubseteq \mathsf{D} \ ,$$
2. yet they respect the strong Bayes vulnerability order, in that
$$\text{for every prior } \pi, \qquad V_1[\pi \triangleright \mathsf{C}] \quad \geq \quad V_1[\pi \triangleright \mathsf{D}] \ ,$$
which we denote by
$$\mathsf{C} \preccurlyeq \mathsf{D} \ .$$

Compositional closure for Dalenius contexts: Justifying refinement - An example

- Example 5 (Continued)

This means that under a very reasonable vulnerability measure (Bayes vulnerability), D never leaks more than C.

With that in mind, we can ask ourselves:

- Are we still justified in insisting that D isn't always safer than C, and claim that $C \not\sqsubseteq D$?
- Equivalently, is it really true that D is sometimes less secure than C?

The answer to both questions is "yes".

We'll now show a convincing Dalenius scenario J that demonstrates that, indeed, D can be less secure than C.

Compositional closure for Dalenius contexts: Justifying refinement - An example

- Example 5 (Continued)

Let the joint-distribution matrix J^{ZX} for J be

J^{ZX}	x_1	x_2	x_3
z_1	$1/3$	$1/3$	0
z_2	0	0	$1/3$

,

so

- its X-marginal is the uniform $(1/3, 1/3, 1/3)$, whereas
- its Z-marginal is $(2/3, 1/3)$.

Given the Dalenius correlation J, what are the Dalenius Bayes vulnerabilities of C and D?

Compositional closure for Dalenius contexts: Justifying refinement - An example

- Example 5 (Continued)

To find them, we use C to carry out the matrix multiplication

J^{ZX}	x_1	x_2	x_3
z_1	$1/3$	$1/3$	0
z_2	0	0	$1/3$

$\cdot$

C	y_1	y_2	y_3
x_1	$1/2$	$1/2$	0
x_2	$1/2$	0	$1/2$
x_3	0	$1/2$	$1/2$

$=$

J^{ZY}	y_1	y_2	y_3
z_1	$1/3$	$1/6$	$1/6$
z_2	0	$1/6$	$1/6$

,

and hence

$$V_1^D(J, \mathrm{C}) \;=\; 1/3 + 1/6 + 1/6 \;=\; 2/3\ .$$

Compositional closure for Dalenius contexts: Justifying refinement - An example

- Example 5 (Continued)

But for D we have

J^{ZX}	x_1	x_2	x_3
z_1	$1/3$	$1/3$	0
z_2	0	0	$1/3$

$\cdot$

D	z_1	z_2
x_1	$2/3$	$1/3$
x_2	$2/3$	$1/3$
x_3	$1/4$	$3/4$

$=$

J^{ZY}	z_1	z_2
z_1	$4/9$	$2/9$
z_2	$1/12$	$1/4$

,

and that gives

$$V_1^D(J,\mathrm{D}) \;=\; 4/9 + 1/4 \;=\; 25/36 \quad > \quad 2/3 \;=\; V_1^D(J,\mathrm{C})\;.$$

Thus the Dalenius Bayes vulnerability with respect to the correlation J is more for D than for C in spite of the fact that $\mathrm{C} \preccurlyeq \mathrm{D}$.

Compositional closure for Dalenius contexts: Justifying refinement - An example

- Example 5 (Continued)

To make the example more concrete, let

- $x_{1,2,3}$ be Jack, John and Jill, and
- $y_{1,2,3}$ be Yes, No and Maybe.

Then C gives the probabilities that each person will give each answer.

E.g., Jack says "Yes" and "No" with equal probability, but never says "Maybe".

C	Yes	No	Maybe
Jack	$1/2$	$1/2$	0
John	$1/2$	0	$1/2$
Jill	0	$1/2$	$1/2$

And no matter what the prior distribution on people might be, the strong Bayes-vulnerability ordering C $\preccurlyeq$ D means that an adversary, having heard what is said, is no more likely with D than with C to guess in one try who it was that spoke.

Compositional closure for Dalenius contexts: Justifying refinement - An example

- Example 5 (Continued)

But now let $z_{1,2}$ be genders, i.e. Male and Female.

Now J expresses the correlation between name and gender as well as the two marginals.

E.g., the person is uniformly likely to be any of Jack, John or Jill, and the gender Male is twice as likely as Female.

J^{ZX}	Jack	John	Jill
Male	$1/3$	$1/3$	0
Female	0	0	$1/3$

Our calculation shows that

$$V_1^D(J, \mathsf{C}) \quad = \quad 2/3 \quad < \quad 25/36 \quad = \quad V_1^D(J, \mathsf{D}) \ ,$$

so an adversary using D is more likely than with C to be able to guess the speaker's gender in one try.

And that is precisely what $\mathsf{C} \not\sqsubseteq \mathsf{D}$ warned us might happen. ◁

Compositional closure for Dalenius contexts: Justifying refinement - An example

- It's important to notice that the construction used in the previous example
  1. is not limited to this particular example, and
  2. is not dependent on the use of concrete channels.
- Indeed, by using Thm. 9.12 (Coriaceous) we conclude that that if

$$C \quad \not\sqsubseteq \quad D$$

then it is guaranteed to exist a Dalenius context J in which

$$V_1^D(J, C) \quad \not\geq \quad V_1^D(J, D) \ ,$$

however "weird" that context might be!

Compositional closure for Dalenius contexts: Justifying refinement - An example

- And indeed it is the case that the Dalenius correlation delivered by Thm. 9.12 (Coriaceous) might look weird.

  That is, it might be difficult to concoct a plausible story of how it could arise in reality.

- But still the conclusion is very striking!

  We must remember that weird contexts are exactly what adversaries devote their time to finding.

- That's why it's important to develop concepts and methods that apply to all contexts.

  We cannot merely focus on the contexts we might think are "reasonable" or "likely".

Bounding Dalenius leakage

Bounding Dalenius leakage

- As we have just seen, Dalenius effects are worrisome.
- As it seems difficult to foresee the correlations that might be discovered to hold between secrets, upper bounds on Dalenius insecurity would be desirable.
- It turns out that Dalenius leakage is a good way to investigate that.

Definition (10.5 – Dalenius g-leakage)

Suppose we are given a channel C with concrete realization $\mathsf{C}\colon \mathcal{X} \to\!\triangleright\, \mathcal{Y}$ and joint distribution $J\colon \mathbb{D}(\mathcal{Z}\times\mathcal{X})$ whose realization J factors into marginal distribution $\rho\colon \mathbb{D}\mathcal{Z}$ and stochastic matrix $\mathsf{B}\colon \mathcal{Z} \to\!\triangleright\, \mathcal{X}$, i.e. so that $\mathsf{J} = \rho \triangleright \mathsf{B}$.

Then, for any gain function $g\colon \mathbb{G}\mathcal{Z}$, define the **Dalenius g-leakage** of J and C by

$$\mathcal{DL}^{\times}_g(J, C) \;:=\; \frac{V_g[\rho \triangleright \mathsf{BC}]}{V_g(\rho)} \;=\; \mathcal{L}^{\times}_g(\rho, \mathsf{BC}) \qquad \text{(multiplicative)}$$

and

$$\mathcal{DL}^{+}_g(J, C) \;:=\; V_g[\rho \triangleright \mathsf{BC}] - V_g(\rho) \;=\; \mathcal{L}^{+}_g(\rho, \mathsf{BC})\,. \qquad \text{(additive)}$$

Bounding Dalenius leakage

- Our first, remarkable result is that:
  - for any channel C and any π and g
  - it is possible to find a correlation J that reduces
    - the multiplicative g-leakage $\mathcal{L}_g^{\times}(\pi, C)$
    - to the multiplicative Dalenius Bayes leakage $\mathcal{DL}_1^{\times}(J, C)$ with respect to that J.

Theorem (10.6)

Let C be a channel on X. For any $\pi : \mathbb{D}\mathcal{X}$ and non-negative gain function $g : \mathbb{G}^+\mathcal{X}$, there exists a correlation J in $\mathbb{D}(\mathcal{Z}\times\mathcal{X})$ such that the multiplicative g-leakage of X is equal to the multiplicative Dalenius Bayes leakage of Z: that is

$$\mathcal{L}_g^{\times}(\pi, C) \quad = \quad \mathcal{DL}_1^{\times}(J, C) \quad .$$

Bounding Dalenius leakage

- **Proof.** We note first that for any ρ we can express $V_g(\rho)$ as $V_g[\rho \triangleright \mathbb{1}]$, where 'the channel $\mathbb{1}$ leaks nothing. (As a matrix, it is a single column of 1's.)

  Then we can use the approach in the proof of Thm. 10.4 (Dalenius necessity) to reason

$$
\begin{array}{lll}
 & \mathcal{L}_g^{\times}(\pi, C) & \\
= & V_g[\pi \triangleright C] / V_g(\pi) & \text{“definition } \mathcal{L}_g^{\times}\text{”} \\
= & V_g[\pi \triangleright C] / V_g[\pi \triangleright \mathbb{1}] & \text{“explained above”} \\
= & c \times V_1^D(J, C) \;/\; c \times V_1^D(J, \mathbb{1}) & \text{“reasoning as in Thm. 10.4,} \\
 & & \text{for } J \text{ depending on } \pi, g\text{”} \\
= & V_1^D(J, C) / V_1^D(J, \mathbb{1}) & \text{“canceling } c\text{'s”} \\
= & V_1[\rho \triangleright \mathsf{BC}] / V_1[\rho \triangleright \mathsf{B}\mathbb{1}] & \text{“use concrete realizations; J} = \rho \triangleright \mathsf{B}\text{”} \\
= & V_1[\rho \triangleright \mathsf{BC}] / V_1(\rho) & \text{“B}\mathbb{1}{=}\mathbb{1}\text{, then again as explained above”} \\
= & \mathcal{DL}_1^{\times}(J, C) \quad , & \text{“Def. 10.5 ”}
\end{array}
$$

  thus establishing the fundamental connection mentioned above between multiplicative Dalenius Bayes leakage under arbitrary correlations and multiplicative g-leakage under arbitrary non-negative gain functions. □

Bounding Dalenius leakage

- With Thm. 10.6 we can prove a very general bound by noticing the following.

  Under Def. 10.5 (Dalenius g-leakage), multiplicative Dalenius g-leakage of C (expressed concretely) is also (ordinary) multiplicative g-leakage of the cascade BC, where as we saw above B is derived from the joint distribution J.

- Expressing multiplicative Dalenius g-leakage as (ordinary) multiplicative g-leakage of a cascade is useful because

  1. we can prove bounds on the multiplicative Bayes capacity of a cascade, and
  2. those bounds carry over (by the Miracle theorem Thm. 7.5) to multiplicative g-leakage for any non-negative gain function g.

Theorem (10.7)

Because we are using cascade, we express this theorem at the concrete level. For any concrete channels $\mathsf{B}\colon \mathcal{Z} \to\!\triangleright \mathcal{X}$ *and* $\mathsf{C}\colon \mathcal{X} \to\!\triangleright \mathcal{Y}$*, we have*

$$\mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{BC}) \quad \leq \quad \min\{\mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{B}),\, \mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{C})\} \quad .$$

Bounding Dalenius leakage

- **Proof.** By the data-processing inequality (Thm. 9.11) we know that for any ρ we have

$$\mathcal{L}_1^{\times}(\rho, \mathsf{BC}) \quad \leq \quad \mathcal{L}_1^{\times}(\rho, \mathsf{B}) \quad \leq \quad \mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{B}) \quad .$$

Hence

$$\mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{BC}) \quad = \quad \sup_{\rho} \mathcal{L}_1^{\times}(\rho, \mathsf{BC}) \quad \leq \quad \mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{B}) \quad .$$

To obtain the upper bound with respect to C, we continue by observing that

$$\begin{array}{lll}
& \mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{BC}) & \\
= & \sum_{y:\mathcal{Y}} \max_{z:\mathcal{Z}} (\mathsf{BC})_{z,y} & \text{“Thm. 7.2”} \\
= & \sum_{y:\mathcal{Y}} \max_{z:\mathcal{Z}} \sum_{x:\mathcal{X}} \mathsf{B}_{z,x}\mathsf{C}_{x,y} & \text{“definition of matrix multiplication”} \\
\leq & \sum_{y:\mathcal{Y}} \max_{z:\mathcal{Z}} \sum_{x:\mathcal{X}} \mathsf{B}_{z,x} \max_{x':\mathcal{X}} \mathsf{C}_{x',y} & \\
= & \sum_{y:\mathcal{Y}} \max_{z:\mathcal{Z}} \max_{x':\mathcal{X}} \mathsf{C}_{x',y} & \text{“row } z \text{ of B sums to 1”} \\
= & \sum_{y:\mathcal{Y}} \max_{x':\mathcal{X}} \mathsf{C}_{x',y} & \\
= & \mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{C}) \quad . & \text{“Thm. 7.2”}
\end{array}$$

Hence we have $\mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{BC}) \leq \min\{\mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{B}), \mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{C})\}$. □

Bounding Dalenius leakage

- Now we can apply Thm. 10.7 above to prove a very general bound on multiplicative Dalenius g-leakage.

Theorem (10.8)

For any channel C, non-negative gain function g, and correlation J, we have $\mathcal{DL}_g^{\times}(J, C) \leq \mathcal{ML}_1^{\times}(\mathbb{D}, C)$.

Proof. Using concrete realizations as usual, we have

$$
\begin{array}{lll}
 & \mathcal{DL}_g^{\times}(\mathsf{J}, \mathsf{C}) & \\
= & \mathcal{L}_g^{\times}(\rho, \mathsf{BC}) & \text{“Def. 10.5”} \\
\leq & \mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{BC}) & \text{“Miracle theorem 7.5”} \\
\leq & \mathcal{ML}_1^{\times}(\mathbb{D}, \mathsf{C}) \quad . & \text{“Thm. 10.7”}
\end{array}
$$

□

- This is an important and robust result, since it does not require us to make any assumptions about the correlation or the gain function (except that the gain function has to be non-negative).