Enhancing Security via Protocol Composition

Catuscia Palamidessi

This talk in a nutshell

- Setting: **Quantitative Information Flow**. Inference attacks using correlation between secret observables
- **Defense**: The system designer can reduce the correlation secret-observables via **protocol composition** (typically randomized)
- **Active Adversary**: may interact with the system and increase the correlation secret-observables
- We formalize the interplay between defender and attacker in **Game Theory**
- Optimal strategy for composition: Saddle points / Nash equilibria. Convex analysis

Quantitative Information Flow

- **General problem:** security and privacy

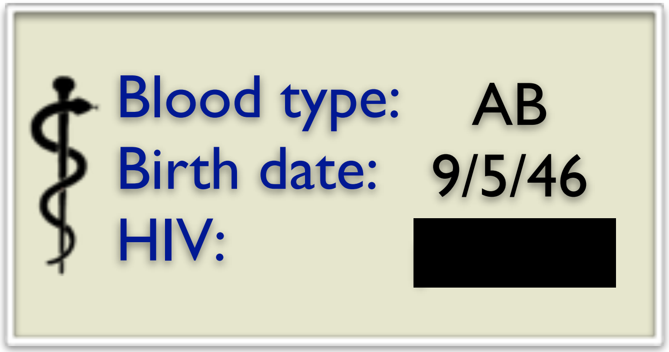

- Access control and encryption are not always sufficient: **systems may leak sensitive information through their correlation with information available to the adversary (observables)**
- Observables: output of the system, public information, side information, physical aspects of the implementation, etc.
- QIF studies measures to assess the threats and techniques to mitigate the leakage due to correlation

Examples of Leakage via correlated observables

Election tabulation
(Sicilian attack)

Execution time
side channel attack

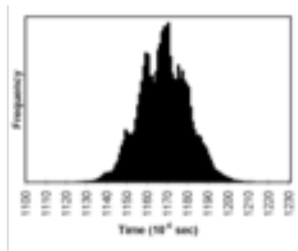

Anonymized
Medical records

Public
voter list

Algorithm
to link
information

De-anonymization
attack

John Smith	HIV

Example

Dining Cryptographers (DC) [Chaum'88]

- A set of nodes with some communication channels (edges).
- One of the nodes (source) wants to broadcast one bit **b** of information
- The source (broadcaster) must remain anonymous

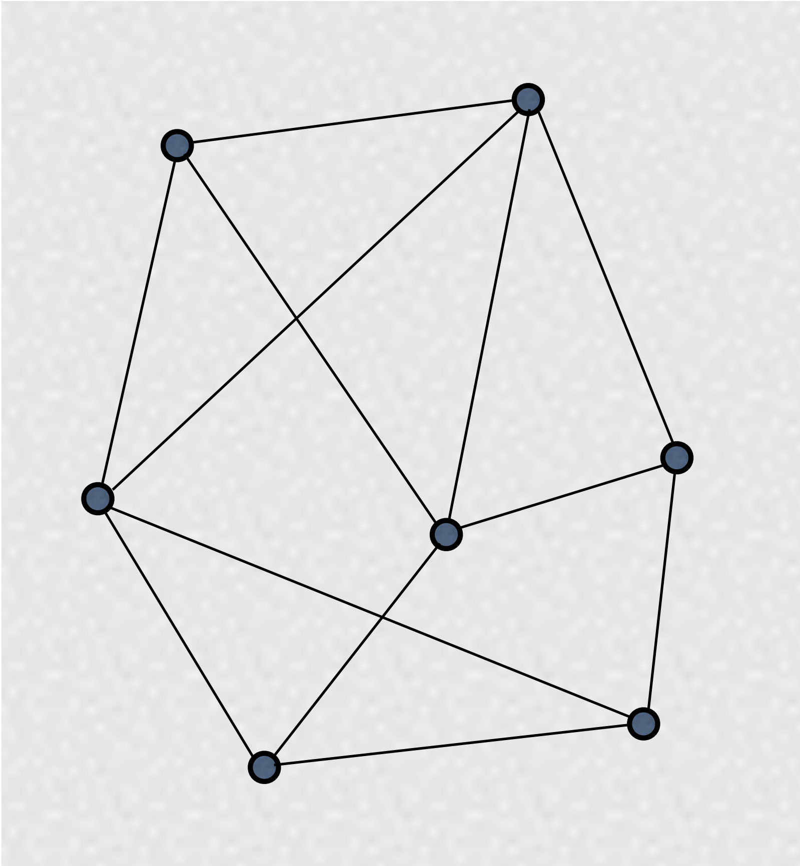

Dining Cryptographers (DC) [Chaum'88]

- A set of nodes with some communication channels (edges).
- One of the nodes (source) wants to broadcast one bit **b** of information
- The source (broadcaster) must remain anonymous

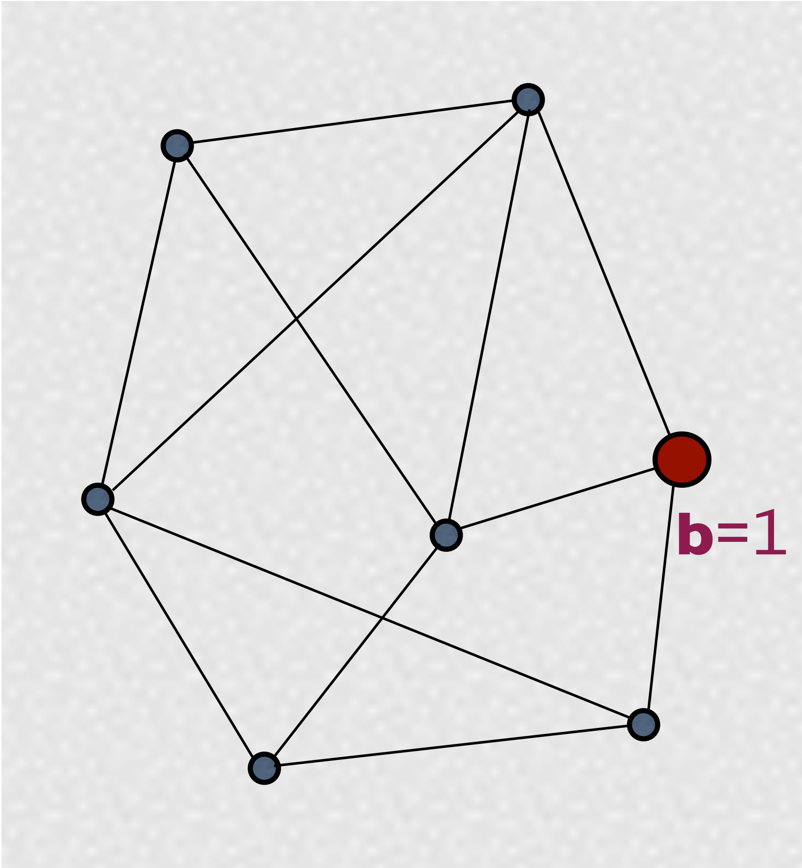

Chaum’s solution

- Associate to each edge a binary coin

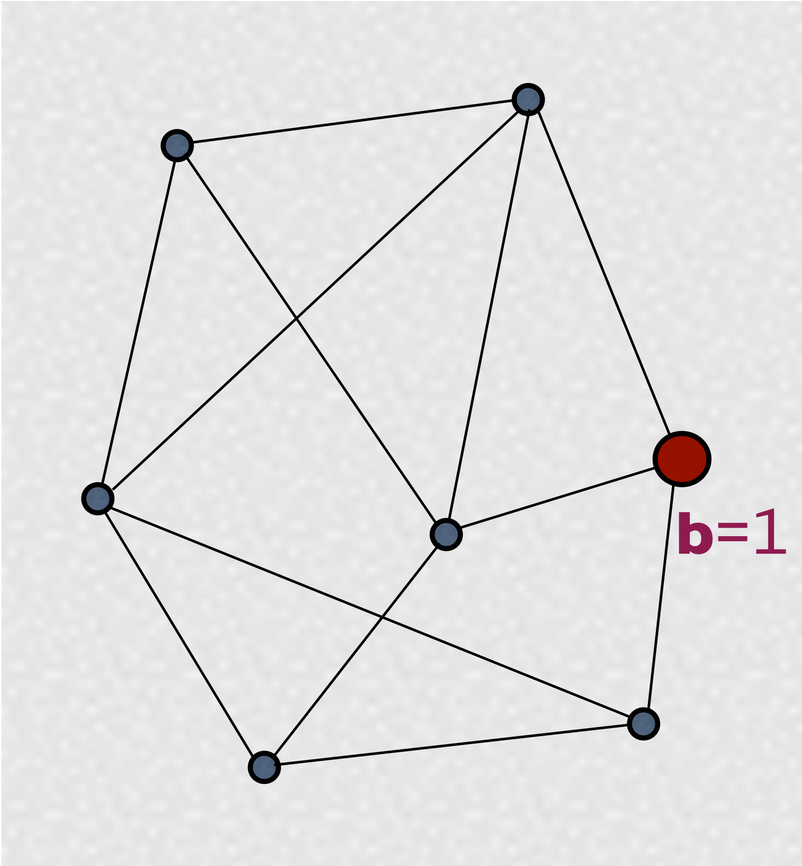

Chaum's solution

- Associate to each edge a binary coin
- Toss the coins

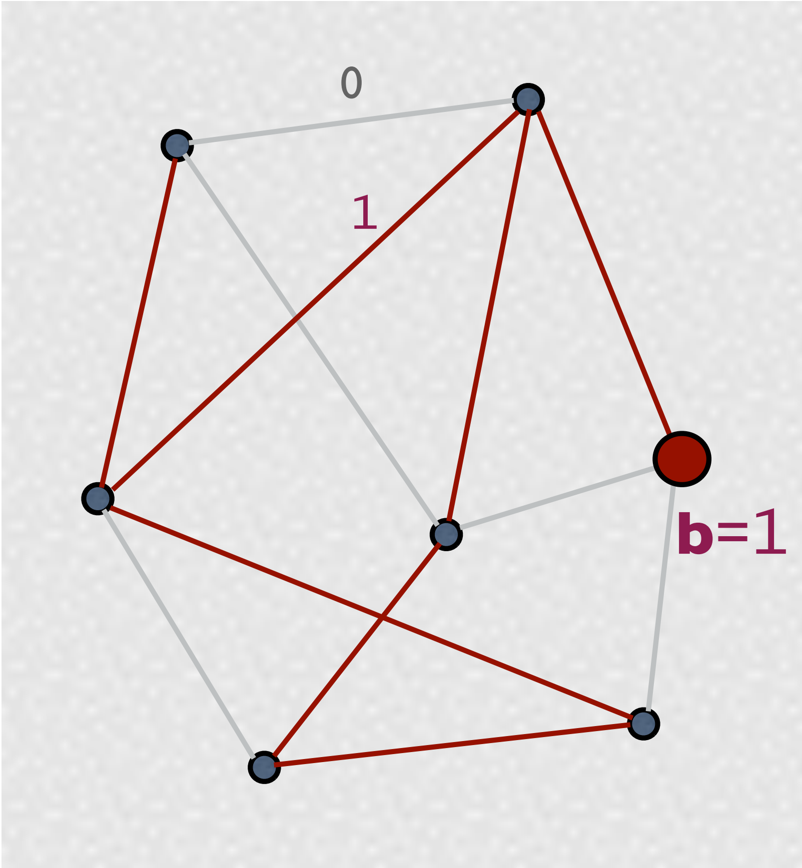

Chaum's solution

- Associate to each edge a binary coin

- Toss the coins

- Each node computes the binary sum of the incident edges. The source adds **b**. They all broadcast their results

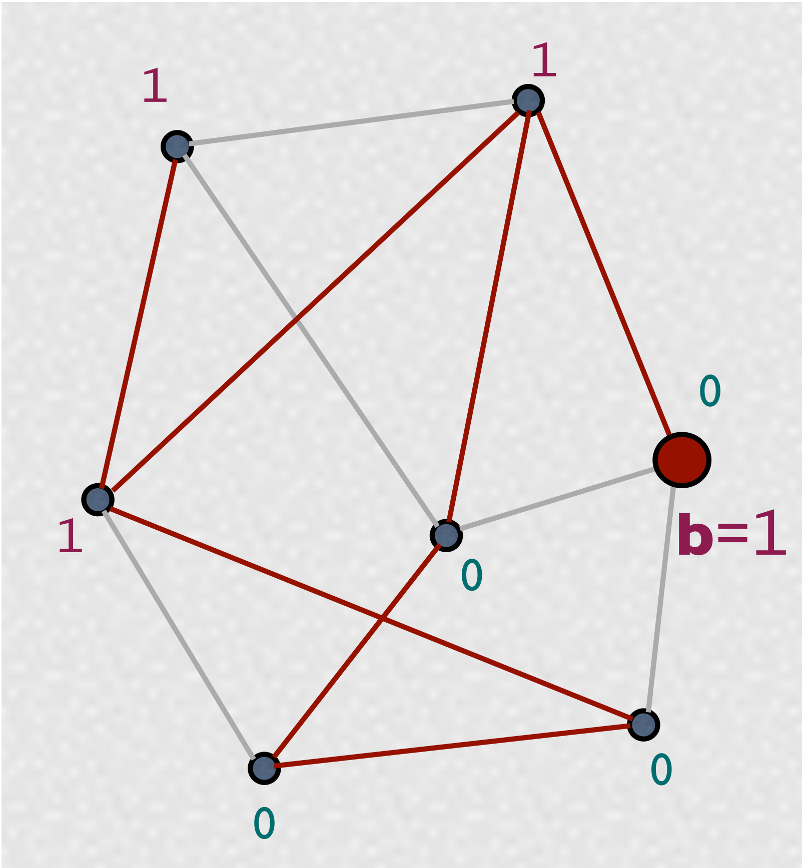

Chaum's solution

- Associate to each edge a binary coin
- Toss the coins
- Each node computes the binary sum of the incident edges. The source adds **b**. They all broadcast their results
- Achievement of the goal: Compute the total binary sum: it coincides with **b**

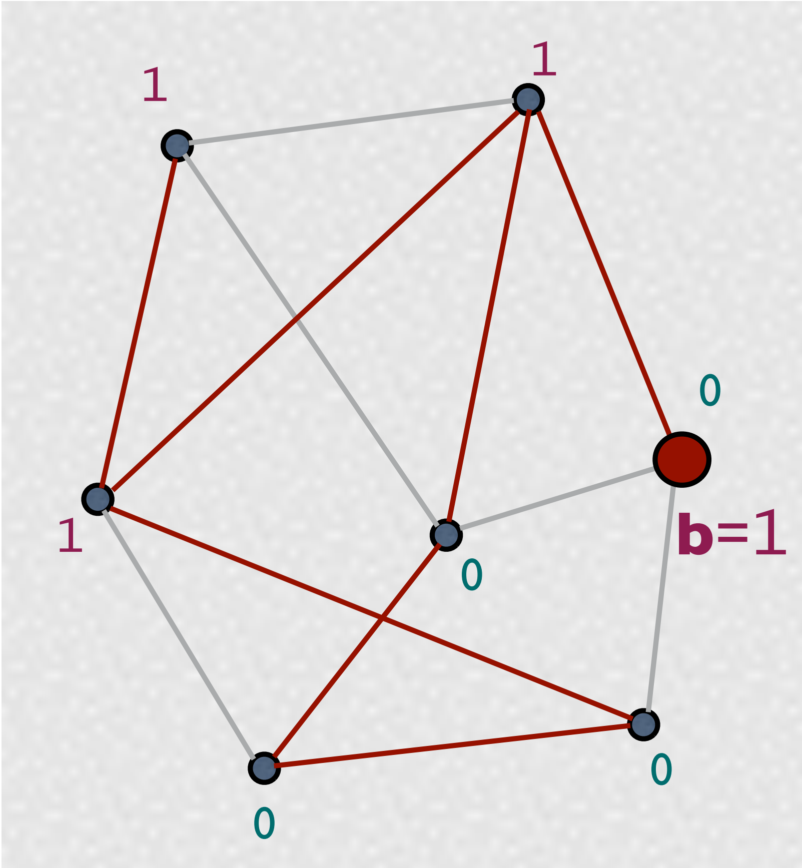

Strong anonymity (Chaum)

- If the graph is connected and the coins are fair, then for an external observer (who observes the declarations of the nodes, but cannot see the value of the coins), the protocol satisfies **strong anonymity**:

> the *a posteriori* probability that a certain node is the source is equal to its *a priori* probability

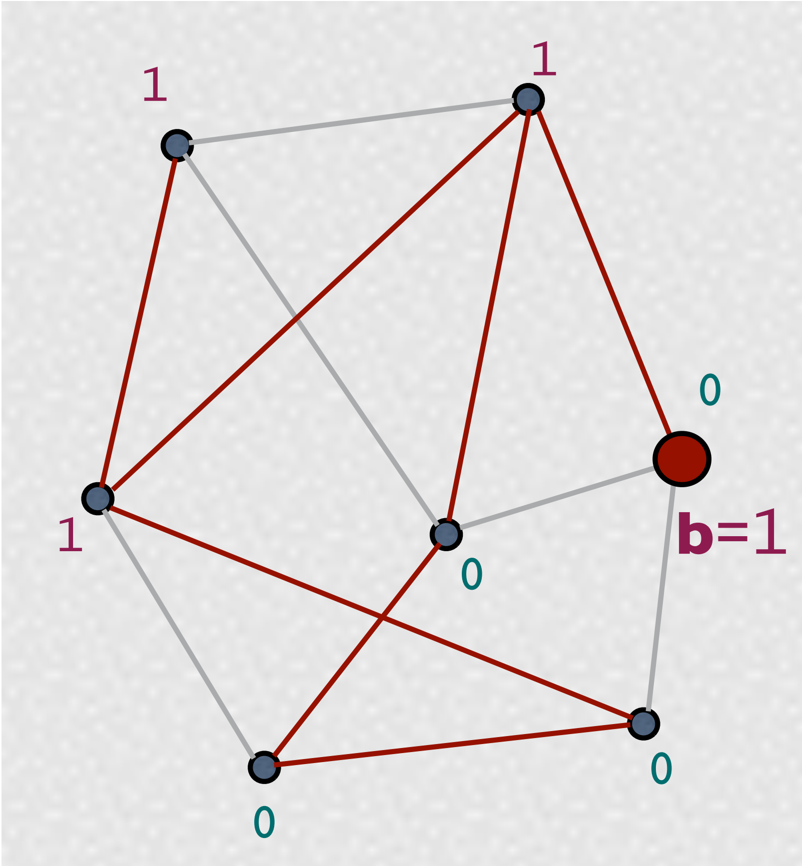

- A priori / a posteriori = before / after observing the declarations

- Note the use of randomization to obfuscate the link between secret and observables

Anonymity of DC Nets

Questions:

What if the coins are biased ?

- Does the protocol still protect the anonymity of the source ? To what extent ?

- How to measure the leakage ?

The basic model:

Systems = Information-Theoretic channels

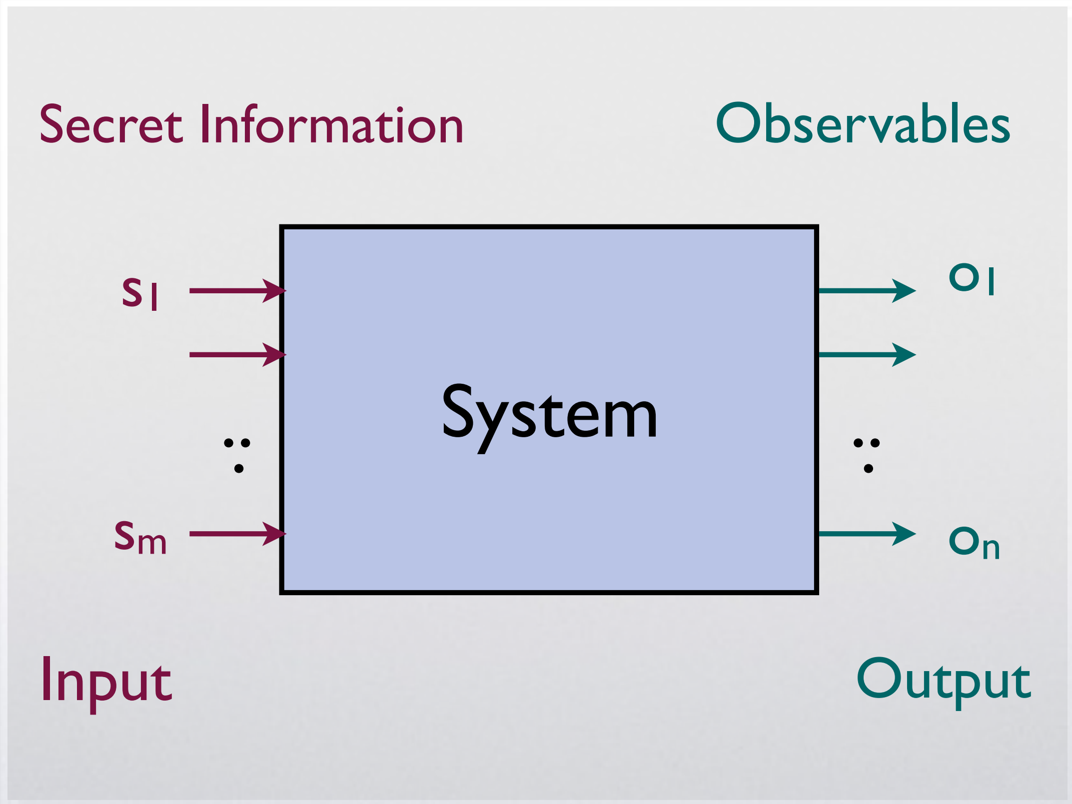

Probabilistic systems are **noisy** channels:
an output can correspond to different inputs, and
an input can generate different outputs, according to a prob. distribution

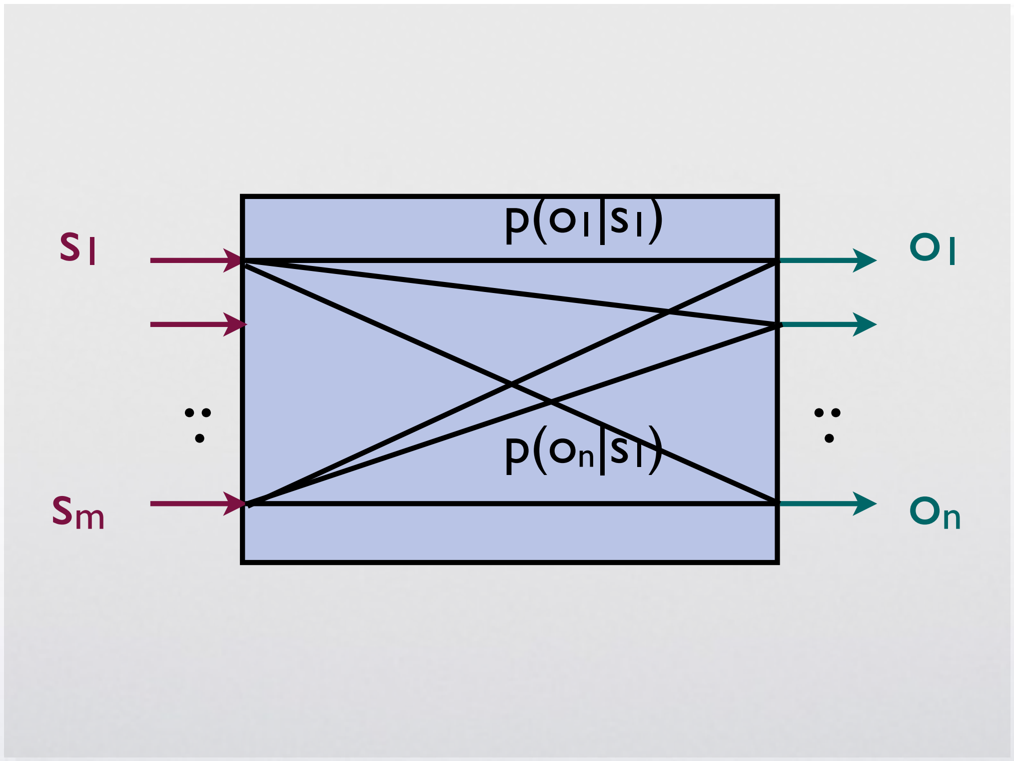

$p(o_j|s_i)$: the conditional probability to observe o_j given the secret s_i

	o_1	$\cdots$	o_n
s_1	$p(o_1 \mid s_1)$	$\cdots$	$p(o_n \mid s_1)$
$\vdots$	$\vdots$		
s_m	$p(o_1 \mid s_m)$		$p(o_n \mid s_m)$

$$p(o|s) = \frac{p(o \ and \ s)}{p(s)}$$

A channel is characterized by its matrix: the array of conditional probabilities

In a information-theoretic channel these conditional probabilities are independent from the input distribution

This means that we can model systems abstracting from the input distribution

Example: DC nets (ring of 3 nodes, b=1)

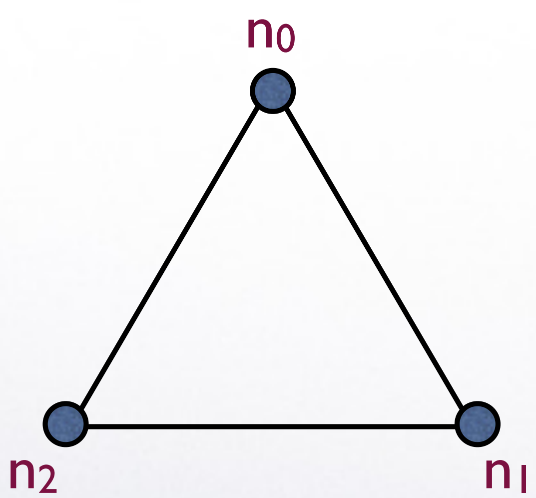

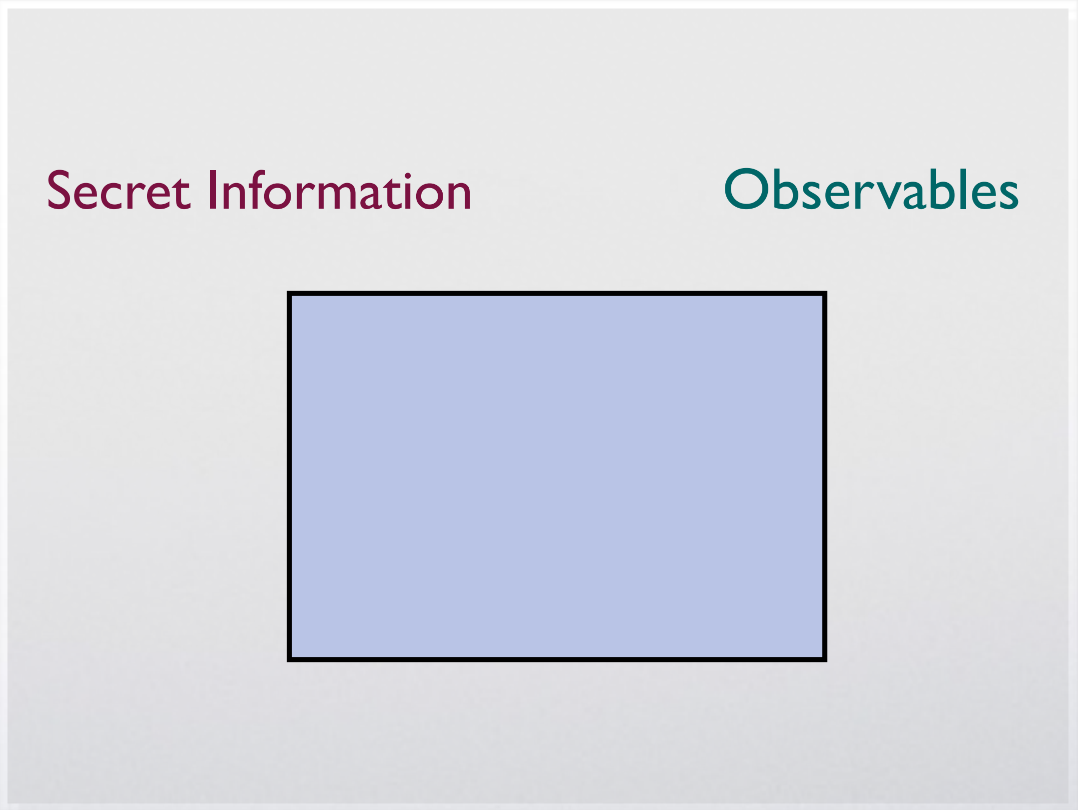

Example: DC nets (ring of 3 nodes, b=1)

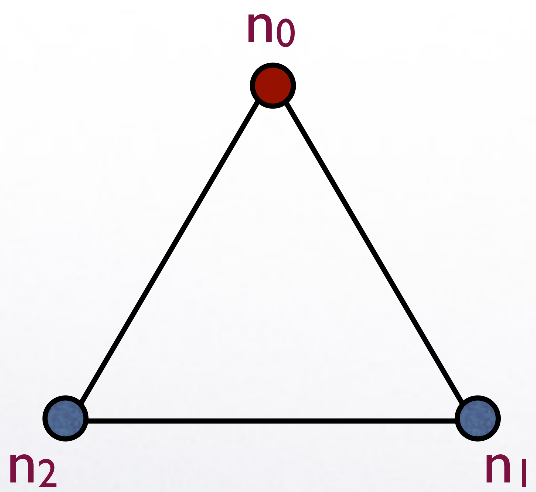

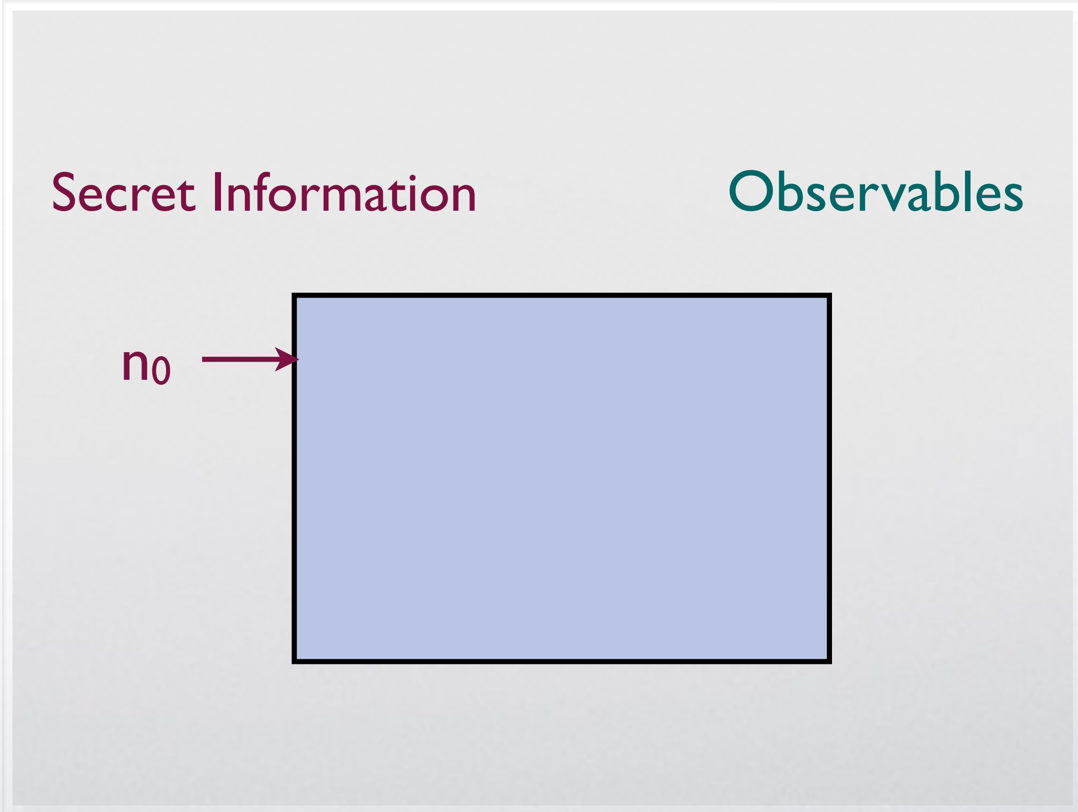

Example: DC nets (ring of 3 nodes, b=1)

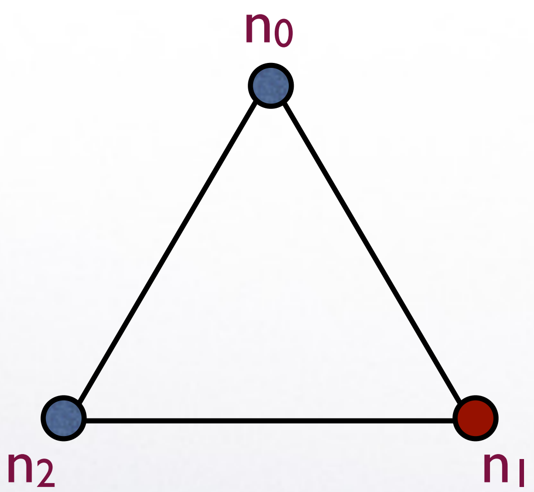

Secret Information

Observables

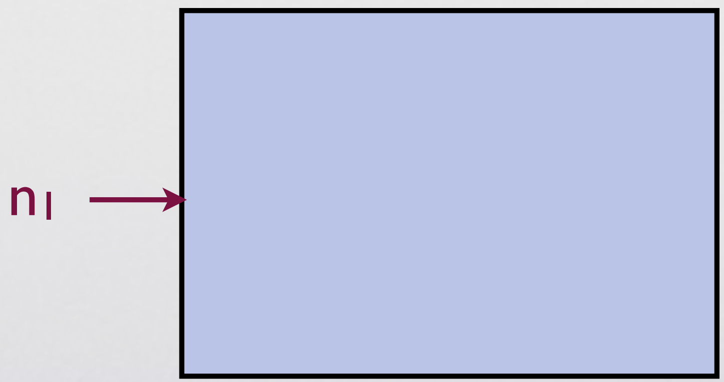

Example: DC nets (ring of 3 nodes, b=1)

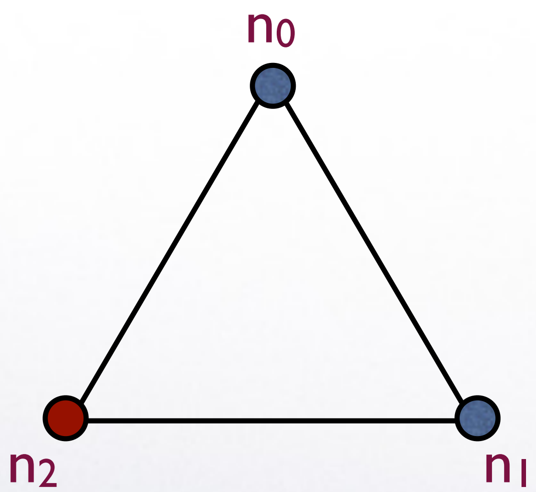

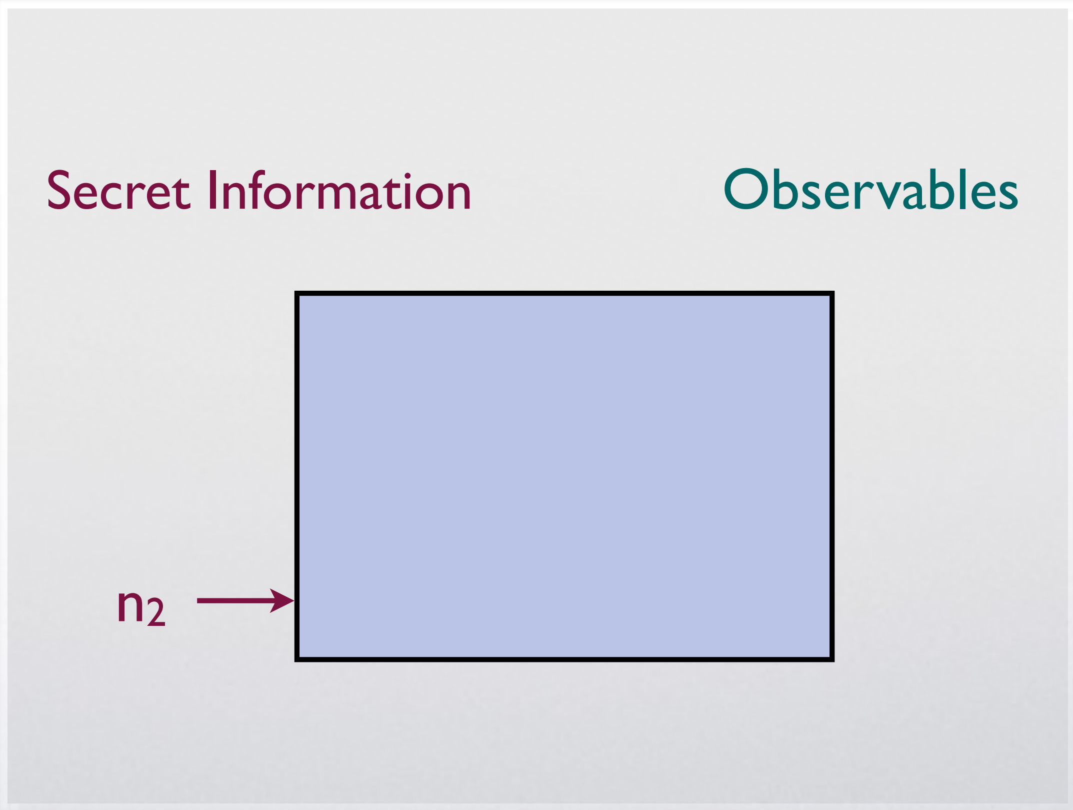

Example: DC nets (ring of 3 nodes, b=1)

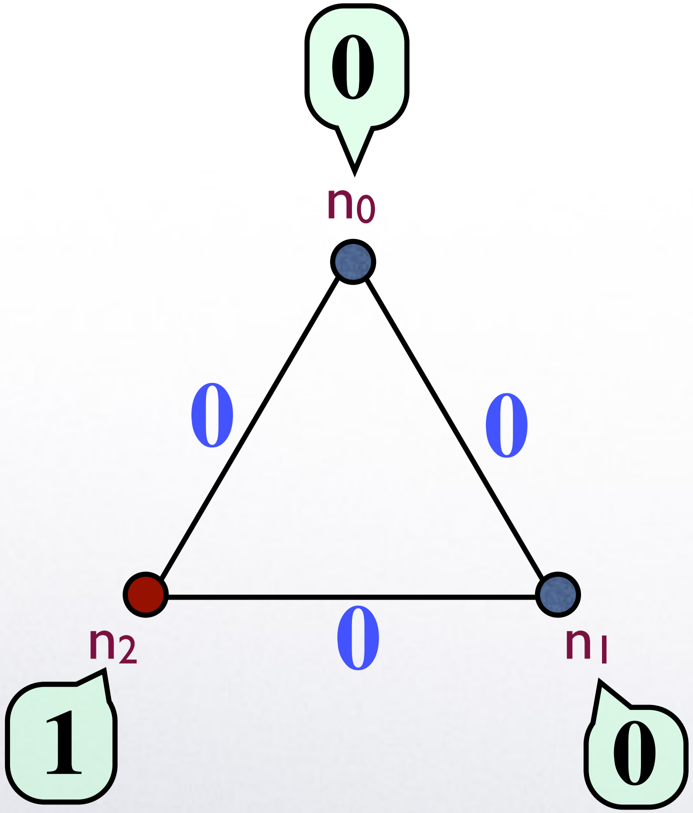

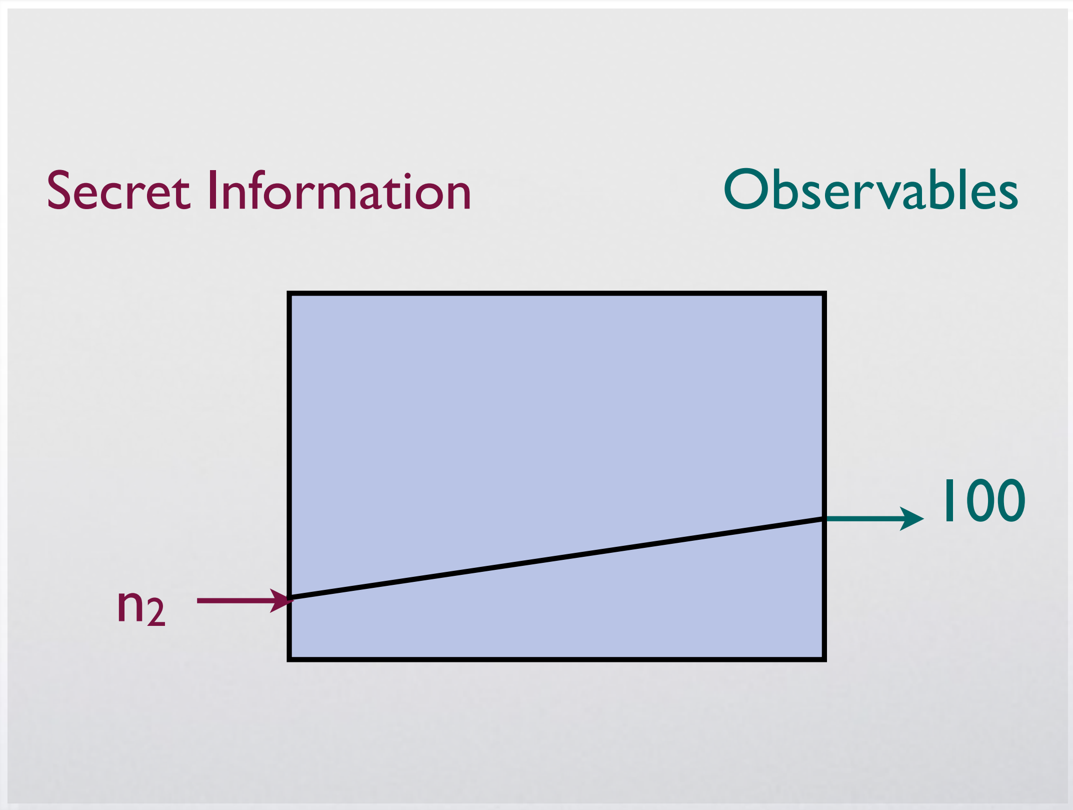

Example: DC nets (ring of 3 nodes, b=1)

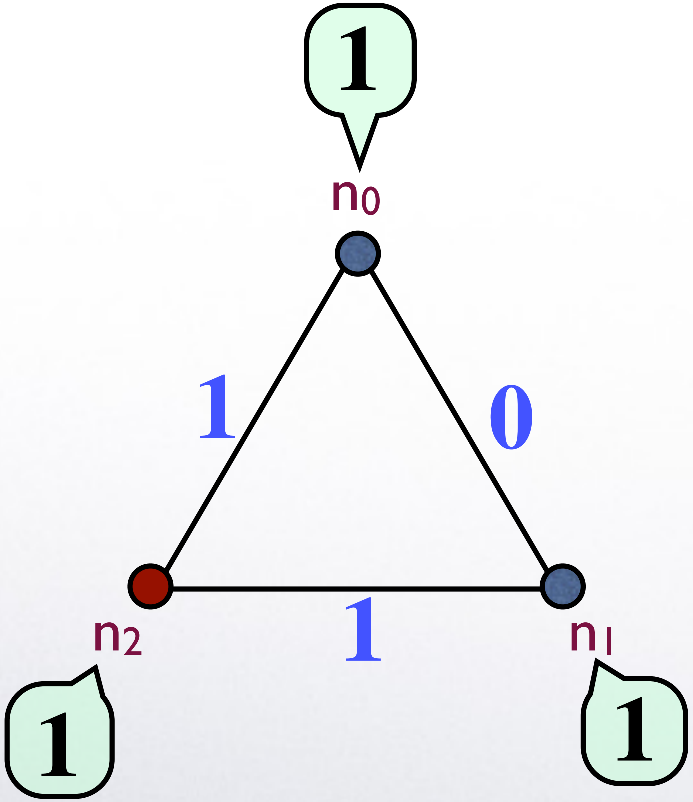

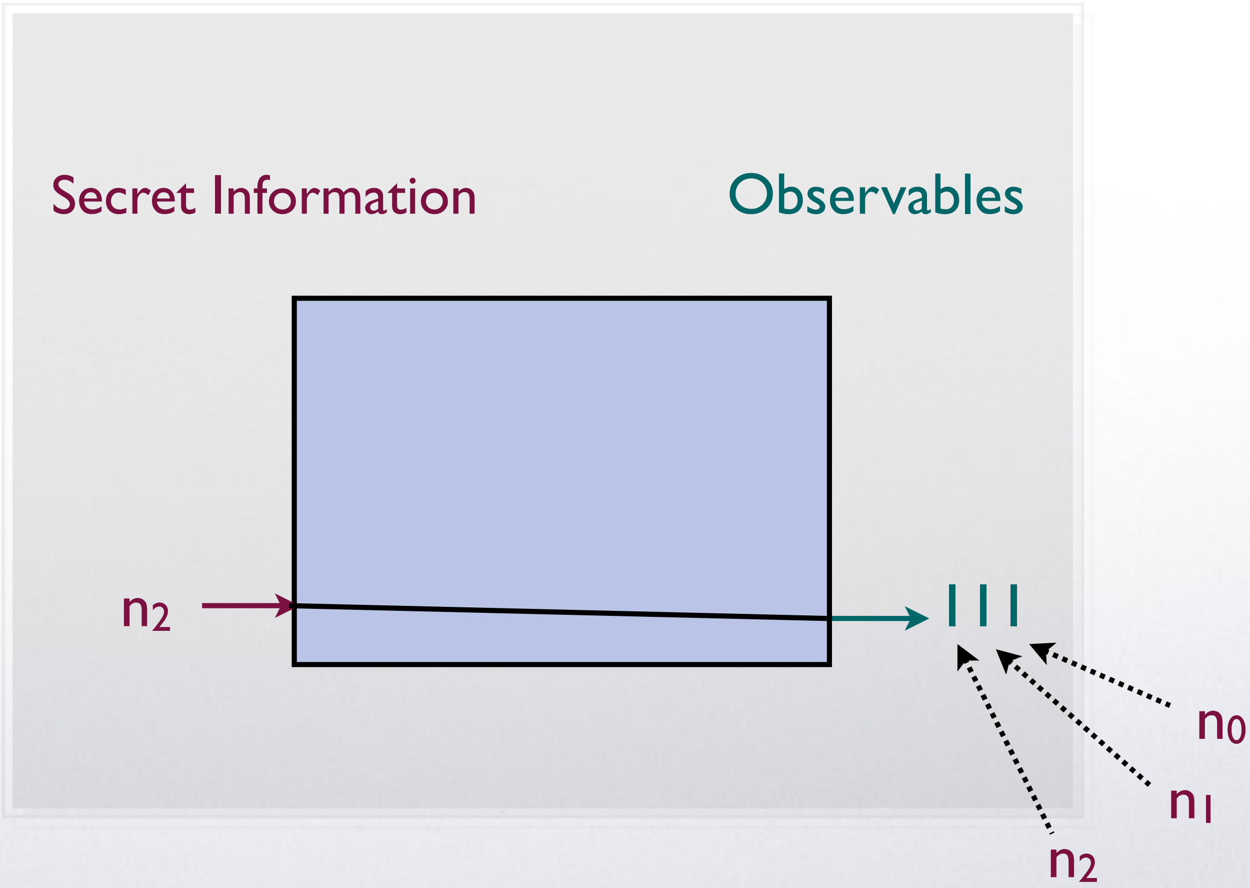

Example: DC nets (ring of 3 nodes, b=1)

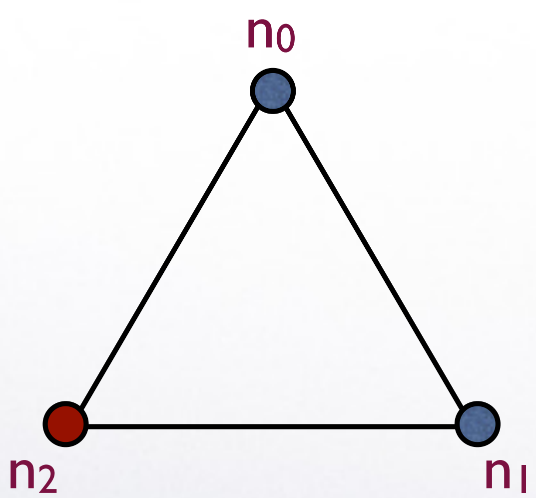

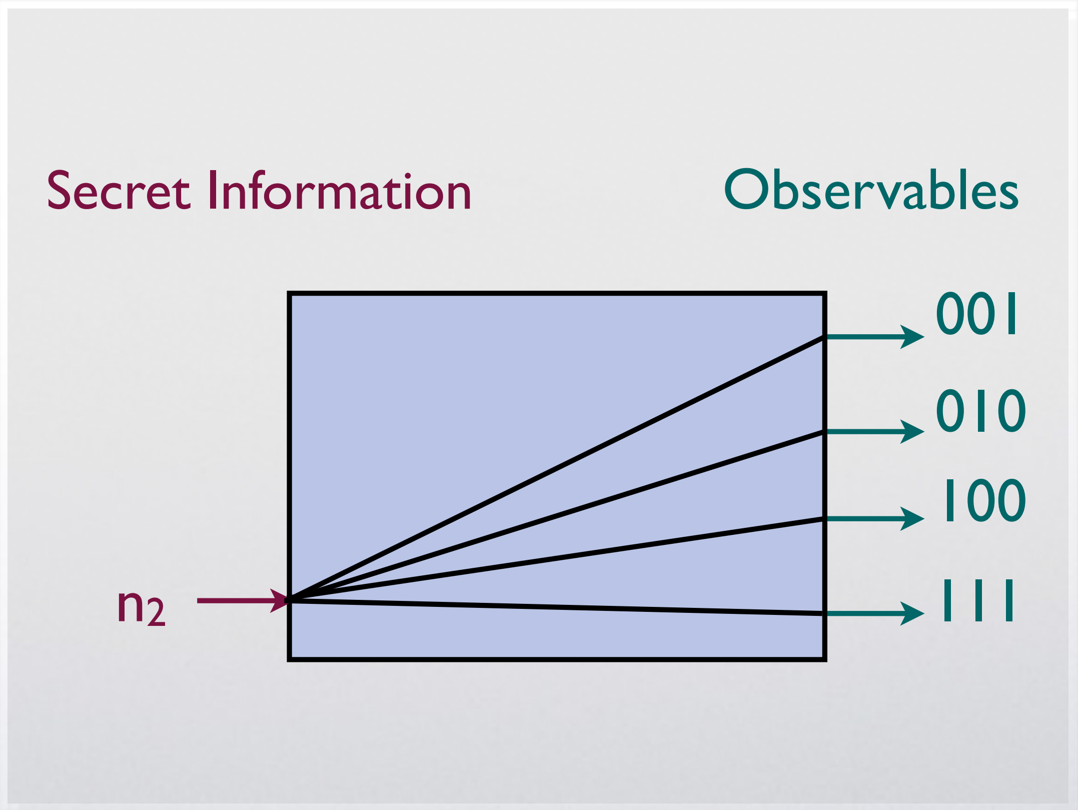

Example: DC nets (ring of 3 nodes, b=1)

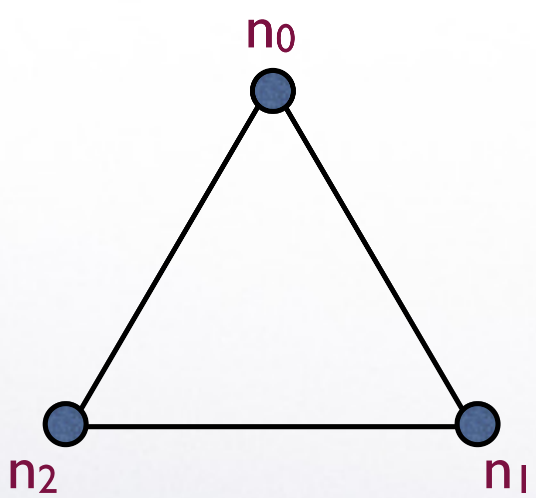

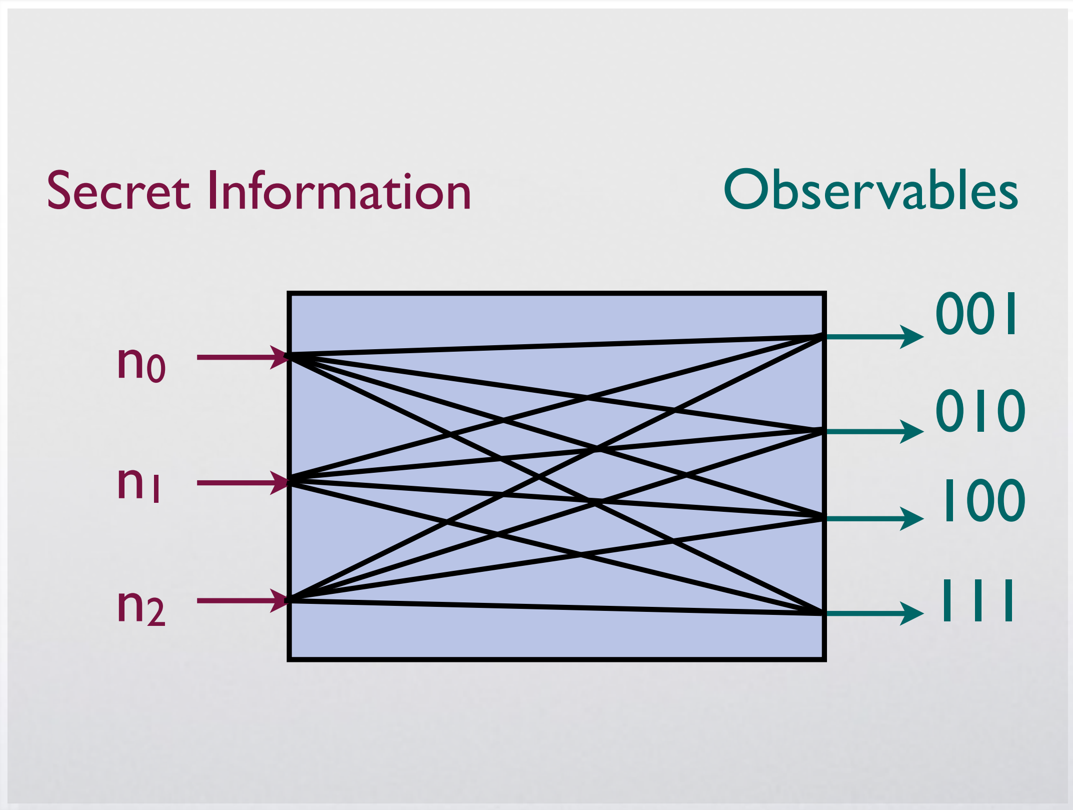

Example: DC nets (ring of 3 nodes, b=1)

	001	010	100	111
n_0	¼	¼	¼	¼
n_1	¼	¼	¼	¼
n_2	¼	¼	¼	¼

	001	010	100	111
n_0	⅓	$^2/_9$	$^2/_9$	$^2/_9$
n_1	$^2/_9$	⅓	$^2/_9$	$^2/_9$
n_2	$^2/_9$	$^2/_9$	⅓	$^2/_9$

fair coins: $Pr(0) = Pr(1) = ½$

strong anonymity

biased coins: $Pr(0) = ⅔$, $Pr(1) = ⅓$

The source is more likely to declare 1 than 0

Quantitative Information Flow

- Intuitively, the **leakage** is the (probabilistic) information that the adversary **gains** about the **secret** through the **observables**

- Each observable **changes** the **prior** probability distribution on the secret values into a **posterior** probability distribution according to the **Bayes** theorem (Bayesian update)

- In average, the information content (about the actual secret value) of the posterior probability is more than or equal to the one of the prior

Bayesian update: prior ⇒ posterior

Bayesian update: prior ⇒ posterior

$\pi(n)$		001	010	100	111
½	n_0	⅓	2⁄9	2⁄9	2⁄9
¼	n_1	2⁄9	⅓	2⁄9	2⁄9
¼	n_2	2⁄9	2⁄9	⅓	2⁄9

prior secret prob

$p(o|n)$
conditional prob

Bayesian update: prior ⇒ posterior

$\pi(n)$
$\frac{1}{2}$
$\frac{1}{4}$
$\frac{1}{4}$

prior secret prob

	001	010	100	111
n_0	$\frac{1}{3}$	$\frac{2}{9}$	$\frac{2}{9}$	$\frac{2}{9}$
n_1	$\frac{2}{9}$	$\frac{1}{3}$	$\frac{2}{9}$	$\frac{2}{9}$
n_2	$\frac{2}{9}$	$\frac{2}{9}$	$\frac{1}{3}$	$\frac{2}{9}$

$p(o|n)$

conditional prob

	001	010	100	111
n_0	$\frac{1}{6}$	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{1}{9}$
n_1	$\frac{1}{18}$	$\frac{1}{12}$	$\frac{1}{18}$	$\frac{1}{18}$
n_2	$\frac{1}{18}$	$\frac{1}{18}$	$\frac{1}{12}$	$\frac{1}{18}$

$p(n,o)$

joint prob

Bayesian update: prior ⇒ posterior

$\pi(n)$	
n_0	½
n_1	¼
n_2	¼

prior secret prob

	001	010	100	111
n_0	⅓	$^2/_9$	$^2/_9$	$^2/_9$
n_1	$^2/_9$	⅓	$^2/_9$	$^2/_9$
n_2	$^2/_9$	$^2/_9$	⅓	$^2/_9$

$p(o|n)$

conditional prob

	001	010	100	111
$p(o)$	$^5/_{18}$	¼	¼	$^2/_9$

obs prob

	001	010	100	111
n_0	⅙	$^1/_9$	$^1/_9$	$^1/_9$
n_1	$^1/_{18}$	$^1/_{12}$	$^1/_{18}$	$^1/_{18}$
n_2	$^1/_{16}$	$^1/_{18}$	$^1/_{12}$	$^1/_{18}$

$p(n,o)$

joint prob

Bayesian update: prior ⇒ posterior

$$p(n|o) = \frac{p(n,o)}{p(o)}$$

p(o)	5⁄18	¼	¼	2⁄9	obs prob

π(n|001)		001	010	100	111
3⁄5	n_0	1⁄3	2⁄9	2⁄9	2⁄9
1⁄5	n_1	2⁄9	1⁄3	2⁄9	2⁄9
1⁄5	n_2	2⁄9	2⁄9	1⁄3	2⁄9
post secret prob					

p(o|n)
conditional prob

	001	010	100	111
n_0	1⁄6	1⁄9	1⁄9	1⁄9
n_1	1⁄18	1⁄12	1⁄18	1⁄18
n_2	1⁄18	1⁄18	1⁄12	1⁄18

p(n,o)
joint prob

If the raws are identical the distribution does not change

If the raws are identical the distribution does not change

$\pi(n)$
½
¼
¼
prior secret prob

	001	010	100	111
n_0	½	¼	⅛	⅛
n_1	½	¼	⅛	⅛
n_2	½	¼	⅛	⅛

p(o|n)
conditional prob

If the raws are identical the distribution does not change

$\pi(n)$
½
¼
¼

prior secret prob

	001	010	100	111
n_0	½	¼	⅛	⅛
n_1	½	¼	⅛	⅛
n_2	½	¼	⅛	⅛

p(o|n)
conditional prob

	001	010	100	111
n_0	¼	⅛	1/16	1/16
n_1	⅛	1/16	1/32	1/32
n_2	⅛	1/16	1/32	1/32

p(n,o)
joint prob

If the raws are identical the distribution does not change

$\pi(n)$		001	010	100	111
½	n_0	½	¼	⅛	⅛
¼	n_1	½	¼	⅛	⅛
¼	n_2	½	¼	⅛	⅛

prior secret prob

$p(o|n)$
conditional prob

	p(o)	½	¼	⅛	⅛
		001	010	100	111
n_0		¼	⅛	1/16	1/16
n_1		⅛	1/16	1/32	1/32
n_2		⅛	1/16	1/32	1/32

obs prob

$p(n,o)$
joint prob

If the raws are identical the distribution does not change

$$p(n \mid o) = \frac{p(o \mid n)\, p(n)}{p(o)}$$

p(o)	½	¼	⅛	⅛	obs prob

π(n)
½
¼
¼

prior secret prob

	001	010	100	111
n_0	½	¼	⅛	⅛
n_1	½	¼	⅛	⅛
n_2	½	¼	⅛	⅛

p(o|n)
conditional prob

	001	010	100	111
n_0	¼	⅛	1/16	1/16
n_1	⅛	1/16	1/32	1/32
n_2	⅛	1/16	1/32	1/32

p(n,o)
joint prob

Formal measures of leakage

Vulnerability $\mathbb{V}(\pi)$ of a secret with prior π

- Represents the "expected damage" the adversary can inflict by making his best guess about the secret value
- The exact definition depends on the operational model of adversary
- Common feature: $\mathbb{V}$ **is convex on** π
- Convexity is a consequence of Data Processing Inequality

Examples:

1. (Converse of) Shannon entropy

   repeated guesses

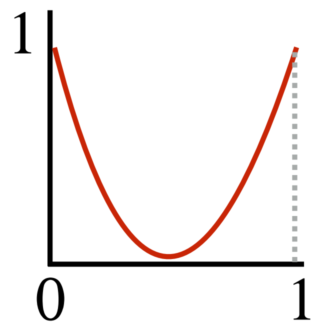

1. Bayes vulnerability: $\mathbb{V}(\pi) = \max_s \pi(s)$

   one guess (one-try attack)

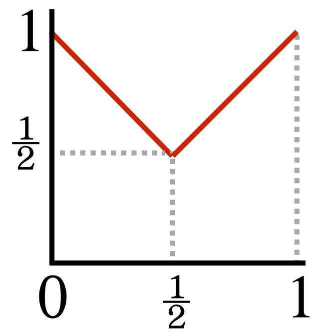

Formal measures of leakage

Posterior Vulnerability $\mathbb{V}(\pi, C)$ of a secret with prior π observed through a channel C

- Vulnerability of the secret after the adversary observes the output: the expected vulnerability of the posterior distributions

$$\mathbb{V}(\pi, C) = \sum_o p(o)\ \mathbb{V}(p(\cdot|o))$$

- Convex in π

Examples

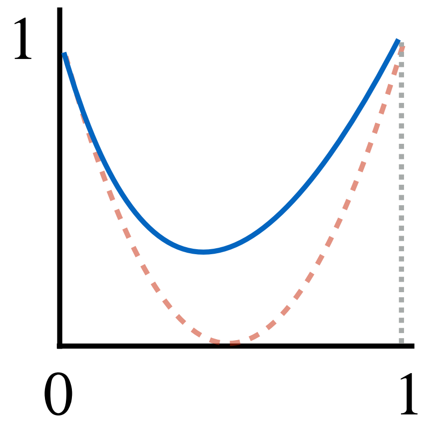

1. Shannon

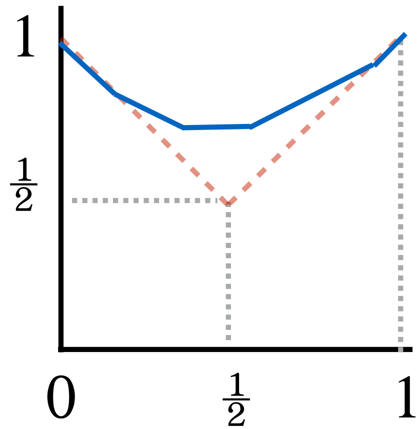

2. Bayes

Convexity

$\mathbb{V}(\pi, C)$ is also convex in C

The defender may lower the vulnerability by randomly combining different channels

Important: random protocols can always be seen as a random combination of deterministic protocols

Convexity

$\mathbb{V}(\pi, C)$ is also convex in C

The defender may lower the vulnerability by randomly combining different channels

Important: random protocols can always be seen as a random combination of deterministic protocols

Example: DC with 2 nodes, 2 biased coins

n_0 n_1

Convexity

$\mathbb{V}(\pi, C)$ is also convex in C

The defender may lower the vulnerability by randomly combining different channels

Important: random protocols can always be seen as a random combination of deterministic protocols

Example: DC with 2 nodes, 2 biased coins

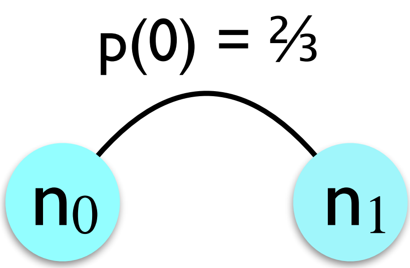

	01	10
n_0	$2/3$	$1/3$
n_1	$1/3$	$2/3$

Convexity

$\mathbb{V}(\pi, C)$ is also convex in C

The defender may lower the vulnerability by randomly combining different channels

Important: random protocols can always be seen as a random combination of deterministic protocols

Example: DC with 2 nodes, 2 biased coins

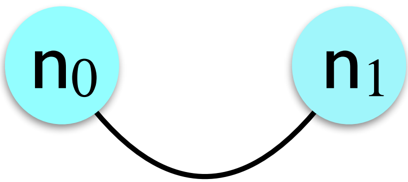

p(1) = ⅔

⅓	⅔
⅔	⅓

Convexity

$\mathbb{V}(\pi, C)$ is also convex in C

The defender may lower the vulnerability by randomly combining different channels

Important: random protocols can always be seen as a random combination of deterministic protocols

Example: DC with 2 nodes, 2 biased coins

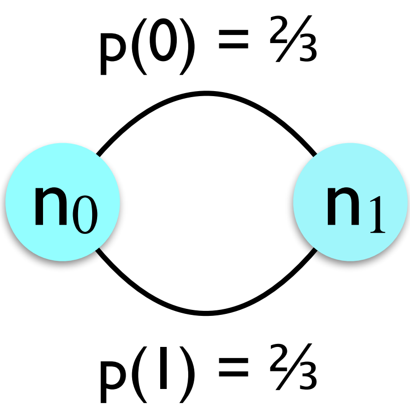

½

⅔	⅓
⅓	⅔

+

½

⅓	⅔
⅔	⅓

=

The completely opaque channel

½	½
½	½

Active adversary

- The adversary may interfere with the system
  - For instance, in the DC, the adversary may control one or more coins
  - More typically, the adversary influences the system by changing the value of some inputs (**low inputs**)
- As a result, the adversary may change the channel matrix as well

Example: the two millionaires problem

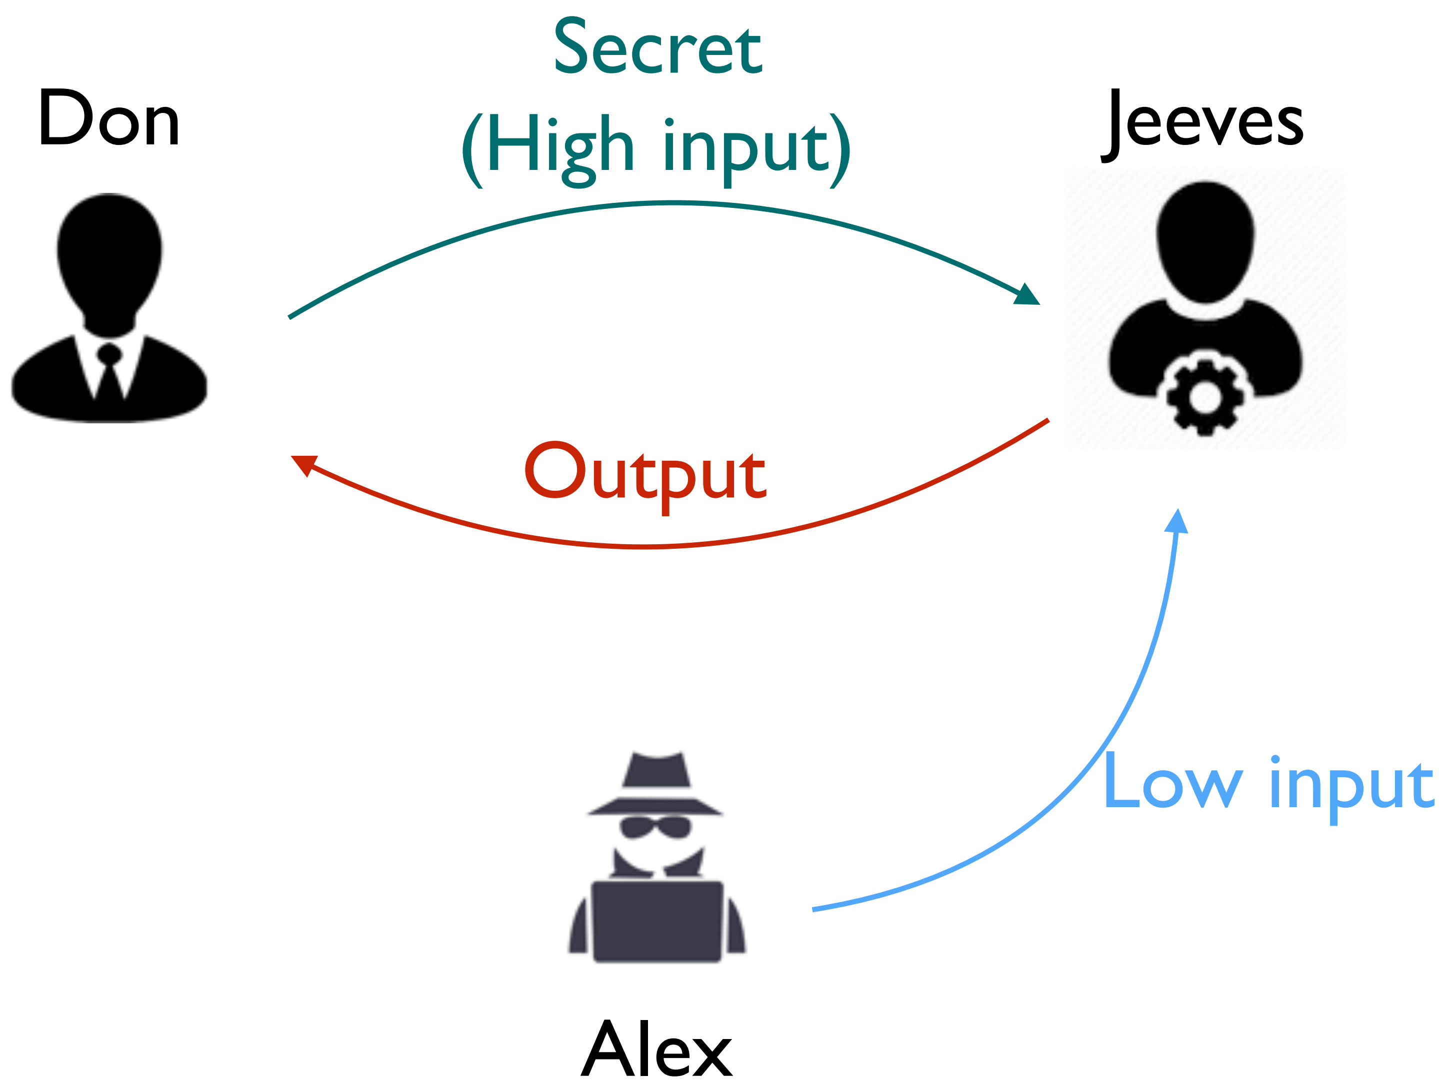

Example: the two millionaires problem

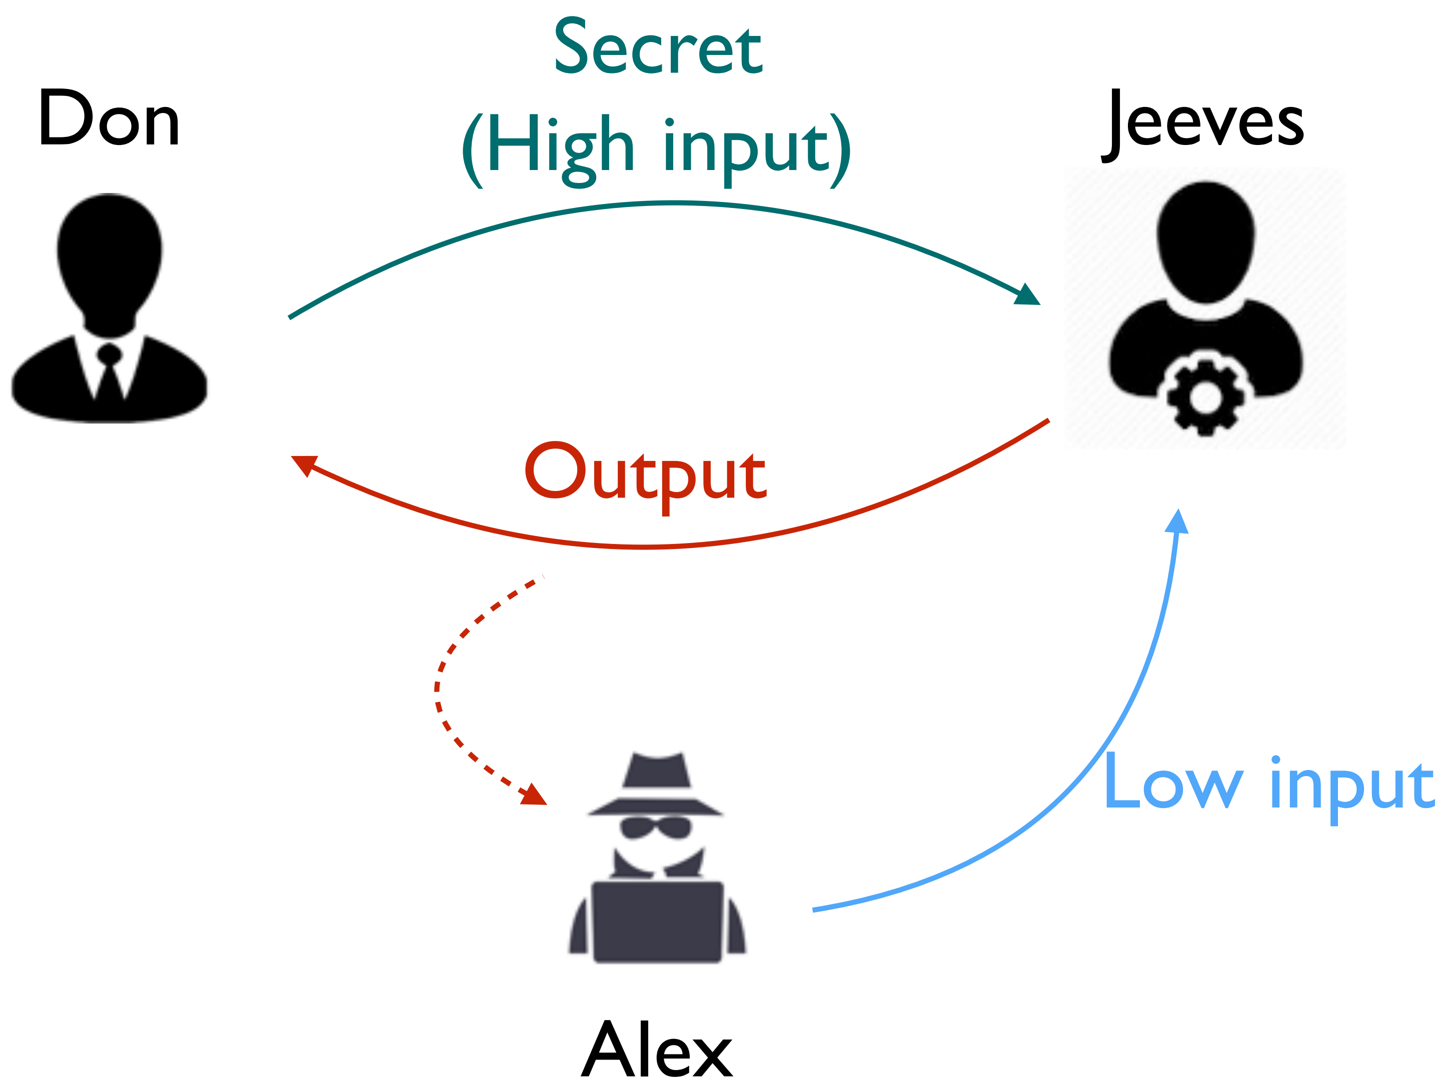

Reducing the vulnerability

Jeeves can run two different programs, both serving the purpose.

Don sends to Jeeves a bit *d* indicating which program he should run

Program 0	Program 1
`High Input:` $x \in \{0,1\}$	`High Input:` $x \in \{0,1\}$
`Low Input:` $a \in \{0,1\}$	`Low Input:` $a \in \{0,1\}$
`Output:` $y \in \{T,F\}$	`Output:` $y \in \{T,F\}$
return $x \leq a$	return $x \geq a$

Depending on the choices a (adversary) and d (defender), we get the following channel matrices:

$a = 0$ / $a = 1$

$d = 0 \quad (x \le a?)$

C_{00}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	0	1

C_{01}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

$d = 1 \quad (x \ge a?)$

C_{10}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

C_{11}	$y = T$	$y = F$
$x = 0$	0	1
$x = 1$	1	0

This can be modeled as a 0-sum game, where the actions a and d are the pure strategies, and the payoff is the leakage (or equivalently, the posterior vulnerability).

The adversary wants to **maximize** the vulnerability, while the defender wants to **minimize** it

Example: Posterior Bayes Vulnerability

$a = 0$

$d = 0 \quad (x \le a?)$

C_{00}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	0	1

$a = 1$

C_{01}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

$d = 1 \quad (x \ge a?)$

C_{10}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

C_{11}	$y = T$	$y = F$
$x = 0$	0	1
$x = 1$	1	0

Example: Posterior Bayes Vulnerability

$a = 0$

$d = 0 \quad (x \le a?)$

C_{00}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	0	1

1

$a = 1$

C_{01}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

$d = 1 \quad (x \ge a?)$

C_{10}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

C_{11}	$y = T$	$y = F$
$x = 0$	0	1
$x = 1$	1	0

1

Example: Posterior Bayes Vulnerability

	$a = 0$				$a = 1$			

$d = 0 \quad (x \leq a?)$

C_{00}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	0	1

1

C_{01}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

½

$d = 1 \quad (x \geq a?)$

C_{10}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

½

C_{11}	$y = T$	$y = F$
$x = 0$	0	1
$x = 1$	1	0

1

Example: Posterior Bayes Vulnerability

$a = 0$ / $a = 1$

$d = 0 \quad (x \leq a?)$

C_{00}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	0	1

1

C_{01}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

½

$d = 1 \quad (x \geq a?)$

C_{10}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

½

C_{11}	$y = T$	$y = F$
$x = 0$	0	1
$x = 1$	1	0

1

Payoff table

$\mathbb{V}$	$a = 0$	$a = 1$
$d = 0$	1	$1/2$
$d = 1$	$1/2$	1

Example: Posterior Bayes Vulnerability

	$a = 0$		$a = 1$	
$d = 0 \quad (x \leq a?)$				

$d = 0 \quad (x \leq a?)$

C_{00}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	0	1

1

C_{01}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

½

$d = 1 \quad (x \geq a?)$

C_{10}	$y = T$	$y = F$
$x = 0$	1	0
$x = 1$	1	0

½

C_{11}	$y = T$	$y = F$
$x = 0$	0	1
$x = 1$	1	0

1

Payoff table

$\mathbb{V}$	$a = 0$	$a = 1$
$d = 0$	1	$1/2$
$d = 1$	$1/2$	1

Similar to the game of the matching pennies

We want to find the optimal strategy (min $\mathbb{V}$) for the defender, taking into account that the adversary will also try to optimize his strategy (max $\mathbb{V}$)

Nash Equilibrium: we have a NE when neither player has any interest to change his strategy unilaterally

We want to find the optimal strategy (min $\mathbb{V}$) for the defender, taking into account that the adversary will also try to optimize his strategy (max $\mathbb{V}$)

Nash Equilibrium: we have a NE when neither player has any interest to change his strategy unilaterally

Usually there is no pure NE,
but there is always a **mixed NE**

Namely, the players can reach a
NE using probabilistic choices

We want to find the optimal strategy (min $\mathbb{V}$) for the defender, taking into account that the adversary will also try to optimize his strategy (max $\mathbb{V}$)

Nash Equilibrium: we have a NE when neither player has any interest to change his strategy unilaterally

Usually there is no pure NE, but there is always a **mixed NE**

Namely, the players can reach a NE using probabilistic choices

$\mathbb{V}$	$a = 0$	$a = 1$
$d = 0$	1	$1/2$
$d = 1$	$1/2$	1

We want to find the optimal strategy (min $\mathbb{V}$) for the defender, taking into account that the adversary will also try to optimize his strategy (max $\mathbb{V}$)

Nash Equilibrium: we have a NE when neither player has any interest to change his strategy unilaterally

Usually there is no pure NE, but there is always a **mixed NE**

Namely, the players can reach a NE using probabilistic choices

	$\mathbb{V}$	$a = 0$	$a = 1$
p	$d = 0$	1	$1/2$
$1-p$	$d = 1$	$1/2$	1

We want to find the optimal strategy (min $\mathbb{V}$) for the defender, taking into account that the adversary will also try to optimize his strategy (max $\mathbb{V}$)

Nash Equilibrium: we have a NE when neither player has any interest to change his strategy unilaterally

Usually there is no pure NE, but there is always a **mixed NE**

Namely, the players can reach a NE using probabilistic choices

		q	$1-q$
	$\mathbb{V}$	$a=0$	$a=1$
p	$d=0$	1	$1/2$
$1-p$	$d=1$	$1/2$	1

We want to find the optimal strategy (min $\mathbb{V}$) for the defender, taking into account that the adversary will also try to optimize his strategy (max $\mathbb{V}$)

Nash Equilibrium: we have a NE when neither player has any interest to change his strategy unilaterally

Usually there is no pure NE, but there is always a **mixed NE**

Namely, the players can reach a NE using probabilistic choices

		q	$1-q$
	$\mathbb{V}$	$a=0$	$a=1$
p	$d=0$	1	$1/2$
$1-p$	$d=1$	$1/2$	1

$$\begin{aligned}\mathbb{V}(p,q) &= 1pq + \tfrac{1}{2}(1-p)q + \tfrac{1}{2}p(1-q) + 1(1-p)(1-q)\\ &= 2pq - p - q + 1\end{aligned}$$

$$\mathbb{V}(p,q) = 2pq - p - q + 1$$

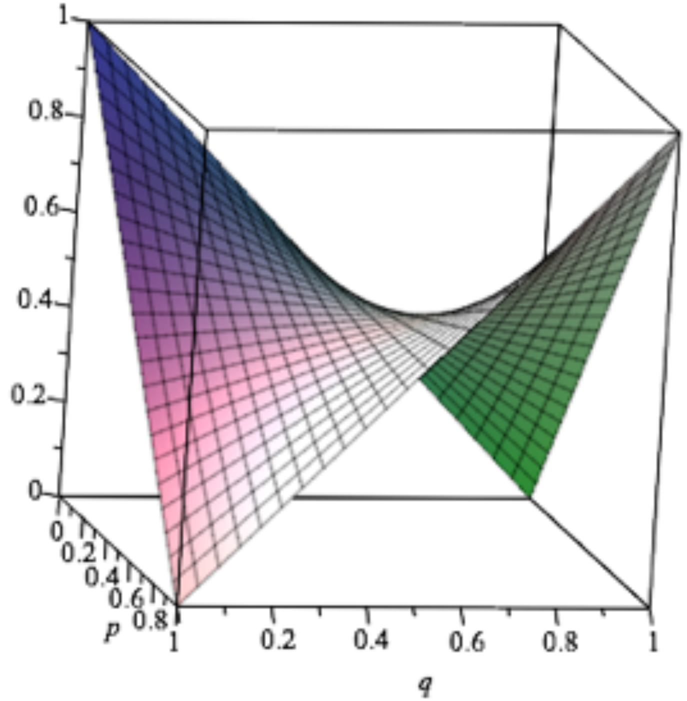

$$\frac{\partial\,\mathbb{V}(p,q)}{\partial p} = \frac{\partial\,\mathbb{V}(p,q)}{\partial q} = 0$$

Van Neumann Th: If $f(p,q)$ is convex in p and concave in q, then

$$\min_p \max_q f(p,q) = \max_q \min_p f(p,q)$$

$$\mathbb{V}(p,q) = 2pq - p - q + 1$$

The NE coincides with the Saddle Point

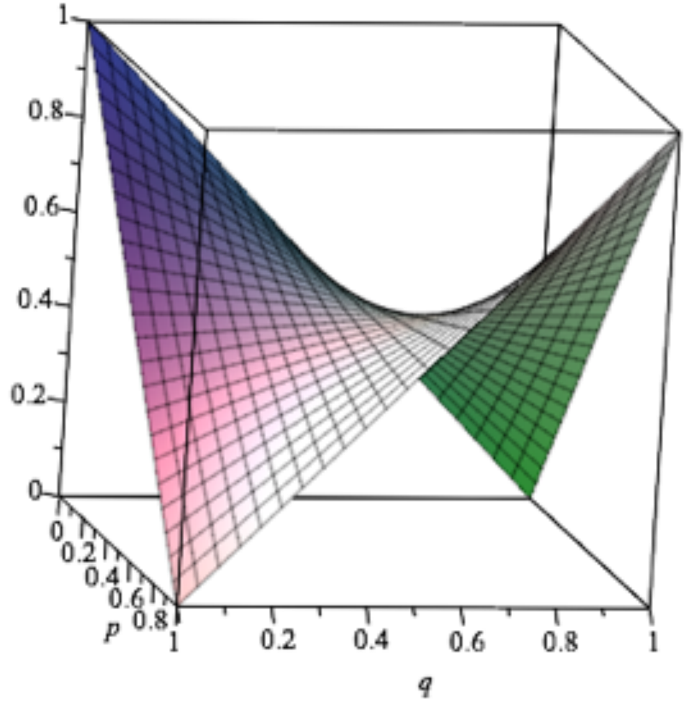

$$\frac{\partial\, \mathbb{V}(p,q)}{\partial p} = \frac{\partial\, \mathbb{V}(p,q)}{\partial q} = 0$$

Van Neumann Th: If $f(p,q)$ is convex in p and concave in q, then

$$\min_p \max_q f(p,q) = \max_q \min_p f(p,q)$$

$$\mathbb{V}(p,q) = 2pq - p - q + 1$$

The NE coincides with the Saddle Point

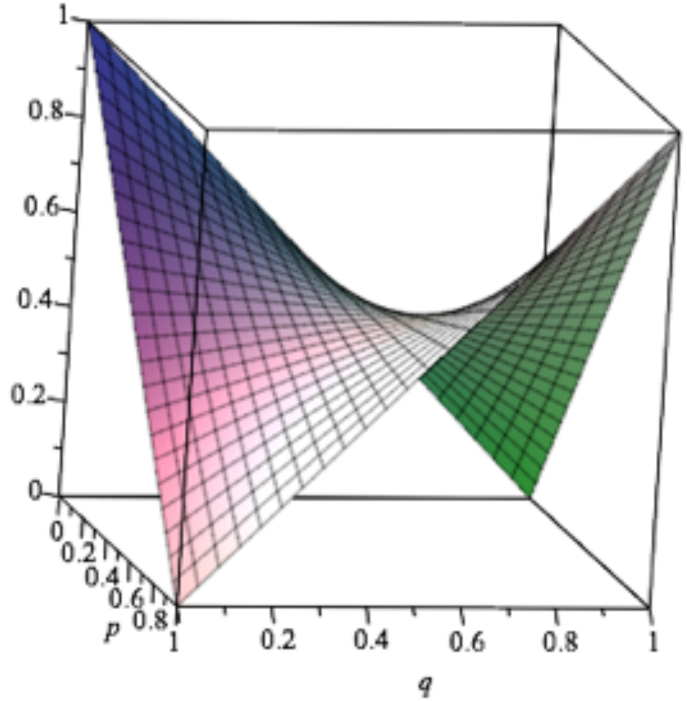

When the partial derivatives exist, the Saddle Point can be computed by imposing $\dfrac{\partial\,\mathbb{V}(p,q)}{\partial p} = \dfrac{\partial\,\mathbb{V}(p,q)}{\partial q} = 0$

Van Neumann Th: If $f(p,q)$ is convex in p and concave in q, then

$$\min_p \max_q f(p,q) = \max_q \min_p f(p,q)$$

$$\mathbb{V}(p,q) = 2pq - p - q + 1$$

The NE coincides with the Saddle Point

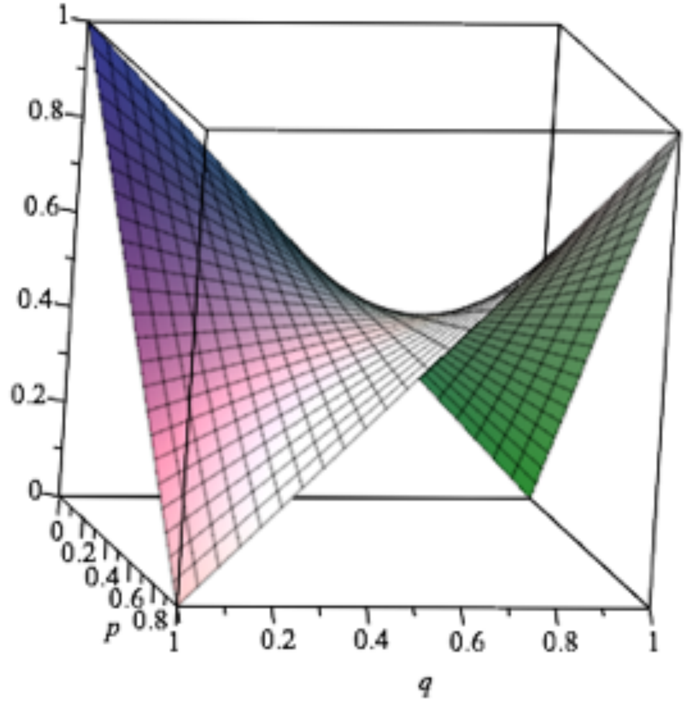

When the partial derivatives exist, the Saddle Point can be computed by imposing $\dfrac{\partial\,\mathbb{V}(p,q)}{\partial p} = \dfrac{\partial\,\mathbb{V}(p,q)}{\partial q} = 0$

In the example, the Saddle Point is for p = q = 1/2

Van Neumann Th: If $f(p,q)$ is convex in p and concave in q, then

$$\min_p \max_q f(p,q) = \max_q \min_p f(p,q)$$

Non-standard games

Assume now that Don wants to know the binary sum

Again, Jeeves has two programs and Don sends to Jeeves a bit *d* indicating which program he should run

Program 0

High Input: $x \in \{0,1\}$
Low Input: $a \in \{0,1\}$
Output: $y \in \{0,1\}$
return $x \oplus a$

Program 1

High Input: $x \in \{0,1\}$
Low Input: $a \in \{0,1\}$
Output: $y \in \{0,1\}$
return $x \oplus a \oplus 1$

Corresponding channel matrices:

$a = 0$

$d = 0 \quad (x \oplus a)$

C_{00}	$y=0$	$y=1$
$x=0$	1	0
$x=1$	0	1

$a = 1$

C_{01}	$y=0$	$y=1$
$x=0$	0	1
$x=1$	1	0

$d = 1 \quad (x \oplus a \oplus 1)$

C_{10}	$y=0$	$y=1$
$x=0$	0	1
$x=1$	1	0

C_{11}	$y=0$	$y=1$
$x=0$	1	0
$x=1$	0	1

Payoff table

$\mathbb{V}$	$a=0$	$a=1$
$d=0$	1	1
$d=1$	1	1

From standard game theory the utility would be 1 and all strategies would be equivalent

However, in the case of our games, the saddle point is (½, ½), and the Utility is 1/2

In fact, if the probability of $d = 0$ is p (strategy of the defender), from the point of view of the adversary the channels are as follows

$a = 0$

C_{p0}	$y = T$	$y = F$
$x = 0$	p	$1 - p$
$x = 1$	$1 - p$	p

$a = 1$

C_{p1}	$y = T$	$y = F$
$x = 0$	$1 - p$	p
$x = 1$	p	$1 - p$

Clearly, the optimal strategy of the defender is for $p = 1/2$, which gives perfectly opaque channels whatever action the attacker chooses

Explanation

- The reason why we get different results than in standard game theory is because the standard utility function is defined as expectation, hence it is affine on the strategies of both players
- In contrast, our games are **convex** on the strategy of the defender (and affine in that of the attacker)
- Van Neumann minimax theorem is still applicable
- Unfortunately, in general the partial derivatives do not exist
- However the saddle point can still be computed by convex analysis

Conclusion

- Probabilistic composition of protocols can be useful to mitigate the Information leakage
- If the attacker is active, then the attacker also has interest to use a probabilistic strategy
- We can model the interplay defender-attacker in Game Theory
- The games are non-standard, but the optimal strategies still exist and can be computed by convex analysis

Future work

- Explore the relation with risk-adverse players
- The goals of the adversary and of the defender may be different => non 0-sum games
- Both adversary and defender may have multiple goals => multiple utility games
- Repeated attacks (repeated runs of the protocol) => repeated games
- Other kinds of interaction => Simultaneous vs alternate games
- Develop the theory of protocol composition (choice and sequential composition)

Thank you !

Questions?