

CS 591: Formal Methods in Security and Privacy

More on Differential Privacy

Marco Gaboardi
gaboardi@bu.edu

Alley Stoughton
stough@bu.edu

Course evaluation

The course evaluation is now available:

<https://bu.campuslabs.com/courseeval/>

Please fill it.

Recording

This is a reminder that we will record the class and we will post the link on Piazza.

This is also a reminder to myself to start recording!

From the previous classes

(ϵ, δ) -Differential Privacy

Definition

Given $\epsilon, \delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ϵ, δ) -differentially private iff
for all adjacent databases D, D' and for every $S \subseteq R$:

$$\Pr[Q(D) \in S] \leq \exp(\epsilon) \Pr[Q(D') \in S] + \delta$$

(ε, δ) -indistinguishability

We can define a ε -skewed version of statistical distance. We call this notion ε -distance.

$$\Delta_\varepsilon(\mu_1, \mu_2) = \sup_{E \subseteq A} \max(\mu_1(E) - e^\varepsilon \mu_2(E), \mu_2(E) - e^\varepsilon \mu_1(E), 0)$$

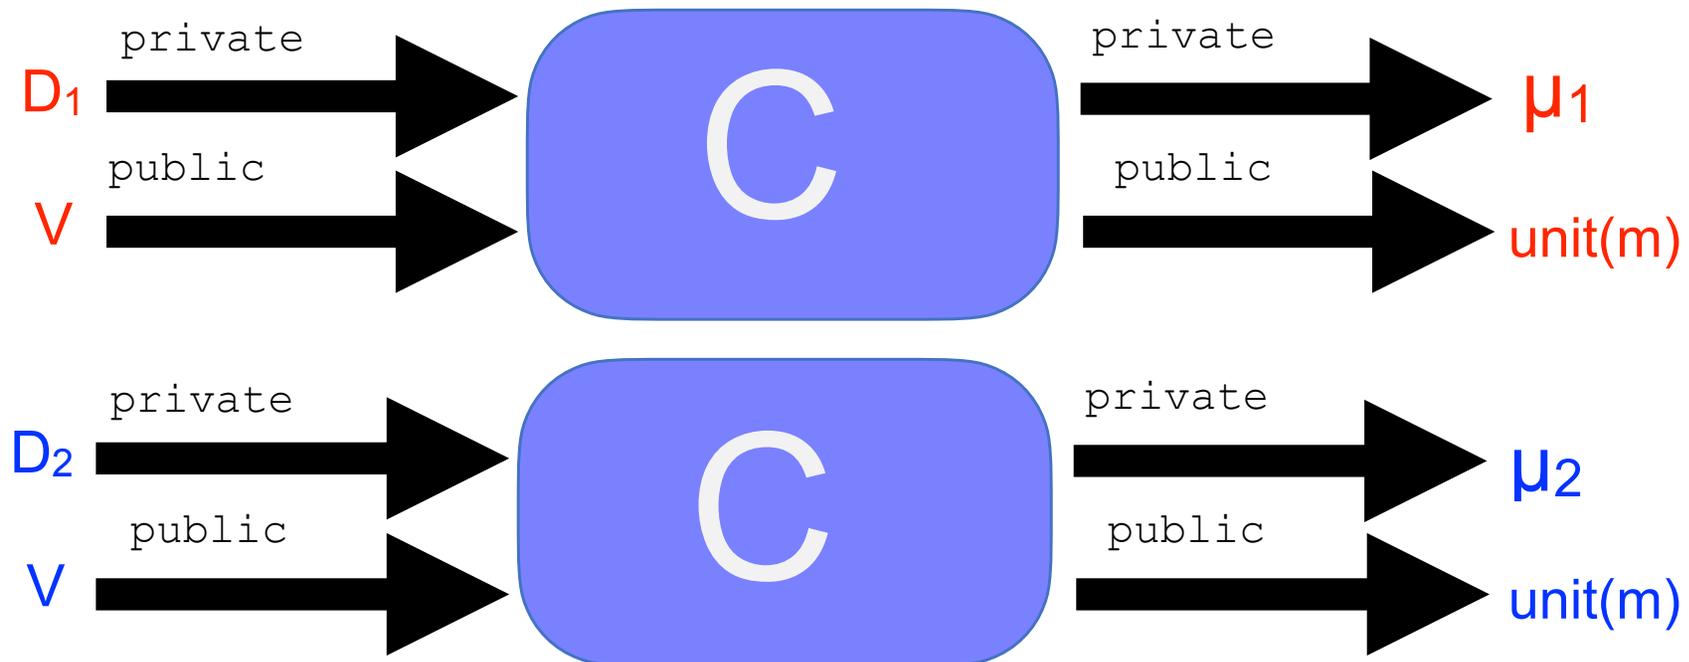
We say that two distributions $\mu_1, \mu_2 \in D(A)$, are at (ε, δ) -indistinguishable if:

$$\Delta_\varepsilon(\mu_1, \mu_2) \leq \delta$$

Differential Privacy as a Relational Property

c is **differentially private** if and only if for every $m_1 \sim m_2$ (extending the notion of adjacency to memories):

$\{c\}_{m_1} = \mu_1$ and $\{c\}_{m_2} = \mu_2$ implies $\Delta_\epsilon(\mu_1, \mu_2) \leq \delta$



apRHL

Indistinguishability
parameter

Precondition
(a logical formula)

$$\vdash_{\epsilon, \delta} C_1 \sim C_2 : P \Rightarrow Q$$

Probabilistic
Program

Probabilistic
Program

Postcondition
(a logical formula)

Validity of apRHL judgments

We say that the quadruple $\vdash_{\varepsilon, \delta} c_1 \sim c_2 : P \Rightarrow Q$ is **valid** if and only if for every pair of memories m_1, m_2 such that $P(m_1, m_2)$ we have:

$\{c_1\}_{m_1} = \mu_1$ and $\{c_2\}_{m_2} = \mu_2$ implies $Q_{\varepsilon, \delta}^*(\mu_1, \mu_2)$.

R- ε - δ -Coupling

Given two distributions $\mu_1 \in D(A)$, and $\mu_2 \in D(B)$, we have an R- ε - δ -coupling between them, for $R \subseteq A \times B$ and $0 \leq \varepsilon$ and $0 \leq \delta \leq 1$, if there are two joint distributions $\mu_L, \mu_R \in D(A \times B)$ such that:

- 1) $\pi_1(\mu_L) = \mu_1$ and $\pi_2(\mu_R) = \mu_2$,
- 2) the support of μ_L and μ_R is contained in R .
That is, if $\mu_L(a, b) > 0$, then $(a, b) \in R$,
and if $\mu_R(a, b) > 0$, then $(a, b) \in R$.
- 3) $\Delta_\varepsilon(\mu_L, \mu_R) \leq \delta$

Today: More on apRHL

Releasing privately the mean of Some Data

```
Mean (d : private data) : public real
  i:=0;
  s:=0;
  while (i<size(d))
    s:=s + d[i];
    i:=i+1;
  z:=$ lap eps s;
  z:= (s/i)+z;
  return z
```

I am using the easycrypt notation here where `lap eps a` corresponds to adding to the value `a` a noise from the Laplace distribution with $b=1/\text{eps}$ and mean $\mu=0$.

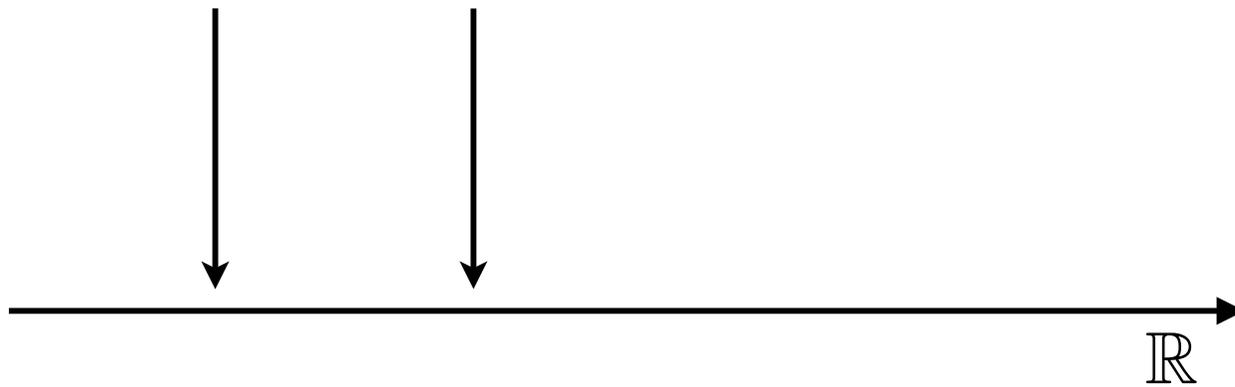
apRHL Laplace

$$\begin{aligned} T_{\varepsilon, 0} & \sim \\ & \begin{aligned} & x_1 := \$ \text{ Lap}(\varepsilon, y_1) \\ & x_2 := \$ \text{ Lap}(\varepsilon, y_2) \\ & : |y_1 - y_2| \leq 1 \quad \Longrightarrow \quad = \end{aligned} \end{aligned}$$

Global Sensitivity

$$GS_q = \max \{ |q(D) - q(D')| \text{ s.t. } D \sim D' \}$$

$q(\text{bu}\{x\})$ $q(\text{bu}\{y\})$



Laplace in EasyCrypt

```
module M = {
  var x: int
  proc f (): int = {
    var r = 0;
    r <$ lap eps x;
    return r;
  }
}.

```

■

```
lemma lem1 : aequiv [ [eps & 0%r]
  M.f ~ M.f
  : ( `|M.x{1} - M.x{2}| <= 1 )
=> res{2} = res{1} ].
proof.
  proc.
  seq 1 1 : ( `|M.x{1} - M.x{2}| <= 1 /\ r{1}=r{2} /\ r{1}=0 ).
  toequiv.
  auto.
(*
  to prove this we can use the lap tactic which takes two
  parameters k1 and k2 and generate two subgoals
  1) k2 * local_eps <= global_eps
  2) |k1 + (M.x{1} - M.x{2})| <= k2
*)
  lap (0) 1.
qed.
```

apRHL

Generalized Laplace

$$x_1 := \$ \text{Lap}(\varepsilon, e_1)$$

$$\vdash_{k_2 * \varepsilon, 0} \sim$$

$$x_2 := \$ \text{Lap}(\varepsilon, e_2)$$

$$\vdash |k_1 + e_1 \langle 1 \rangle - y_2 \langle 2 \rangle| \leq k_2$$

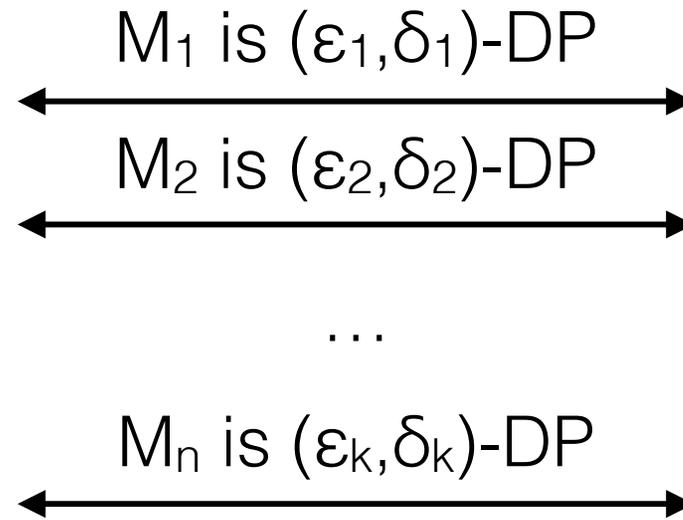
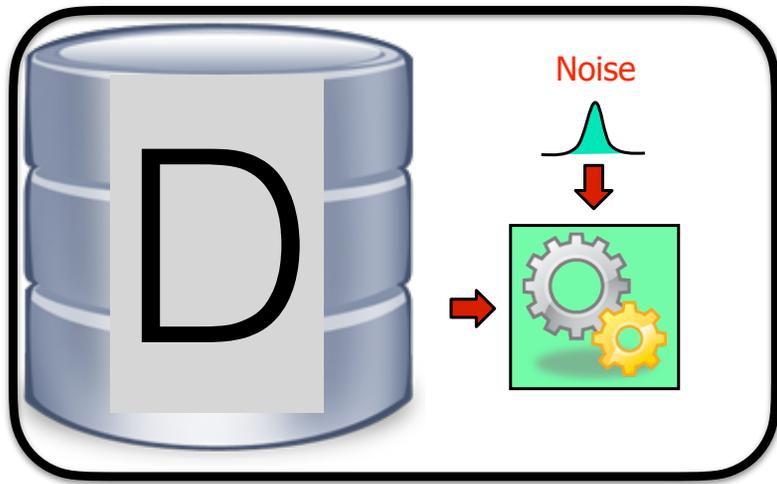
$$\implies x_1 \langle 1 \rangle + k_1 = x \langle 2 \rangle$$

Releasing partial sums

```
DummySum (d : {0,1} list) : real list
  i := 0;
  s := 0;
  r := [];
  while (i < size d)
    s := s + d[i];
    z := $ lap eps s;
    r := r ++ [z];
    i := i + 1;
  return r
```

I am using the easycrypt notation here where $\text{Lap}(\text{eps}, a)$ corresponds to adding to the value a a noise from the Laplace distribution with $b=1/\text{eps}$ and mean $\mu=0$.

Composition



The overall process is $(\epsilon_1 + \epsilon_2 + \dots + \epsilon_k, \delta_1 + \delta_2 + \dots + \delta_k)$ -DP

apRHL Composition

$$\vdash_{\varepsilon_1, \delta_1} C_1 \sim C_2 : P \Rightarrow R \quad \vdash_{\varepsilon_2, \delta_2} C_1' \sim C_2' : R \Rightarrow S$$

$$\vdash_{\varepsilon_1 + \varepsilon_2, \delta_1 + \delta_2} C_1 ; C_1' \sim C_2 ; C_2' : P \Rightarrow S$$

This corresponds to EC command:

$\text{seq } 1 \ 1 : (\text{postcondition } R) \langle [\text{eps}_1 \ \& \ \text{del}_1] \rangle.$

Releasing partial sums

```
DummySum (d : {0,1} list) : real list
  i := 0;
  s := 0;
  r := [];
  while (i < size d)
    s := s + d[i];
    z := $ lap eps s;
    r := r ++ [z];
    i := i + 1;
  return r
```

apRHL awhile

$$P \wedge e \langle 1 \rangle \leq 0 \implies \neg b \langle 1 \rangle$$

$$\begin{aligned} & \vdash \varepsilon_k, \delta_k \quad c1 \sim c2 : P \wedge b \langle 1 \rangle \wedge b \langle 2 \rangle \wedge k = e \langle 1 \rangle \\ & \quad \wedge e \langle 1 \rangle \leq n \\ & \implies P \wedge b \langle 1 \rangle = b \langle 2 \rangle \wedge k < e \langle 1 \rangle \end{aligned}$$

$$\begin{aligned} & \vdash \sum \varepsilon_k, \sum \delta_k \quad \text{while } b1 \text{ do } c1 \sim \text{while } b2 \text{ do } c2 \\ & \quad : P \wedge b \langle 1 \rangle = b \langle 2 \rangle \wedge e \langle 1 \rangle \leq n \\ & \implies P \wedge \neg b \langle 1 \rangle \wedge \neg b \langle 2 \rangle \end{aligned}$$

Parallel Composition

Let $M_1:DB \rightarrow R$ be a (ϵ_1, δ_1) -differentially private program and $M_2:DB \rightarrow R$ be a (ϵ_2, δ_2) -differentially private program. Suppose that we partition D in a data-independent way into two datasets D_1 and D_2 . Then, the composition $M_{1,2}:DB \rightarrow R$ defined as

$$MP_{1,2}(D) = (M_1(D_1), M_2(D_2))$$

is $(\max(\epsilon_1, \epsilon_2), \max(\delta_1, \delta_2))$ -differentially private.

Parallel Composition In EC

```
lemma lem2par : aequiv [ [eps & del]
  M1.f ~ M1.f
  : ( `|M1.d{1}.`1 - M1.d{2}.`1| <= 1 /\ M1.d{1}.`2 = M1.d{2}.`2)
=> res{2} = res{1}].
```

```
proof.
```

```
  proc.
```

```
  seq 3 3 : ( `|z{1} - z{2}| <= 1 /\ y{1} = y{2} ).
```

```
  wp. toequiv. skip. trivial.
```

```
  seq 1 1 : (r{1}=r{2} /\ y{1} = y{2}) <[ eps & del ]>.
```

```
  lap (0) 1 => //.
```

```
  lap (0) 0.
```

```
(*
```

we could have just avoid using laplace at all in the algorithm,
but the point of this example is to show how the same algorithm
can be analyzed in different ways. We will also see that this idea
is behind the idea of parallel composition

```
*)
```

```
qed.
```

Releasing partial sums

```
DummySum (d : {0,1} list) : real list
  i:=0;
  s:=0;
  r:=[];
  while (i<size d)
    z:=$ lap eps d[i];
    s:= s + z;
    r:= r ++ [s];
    i:= i+1;
  return r
```

Parallel Composition In EC

```
Lemma noisy_sum n j : 0 <= j < n => aequiv [ [ eps & 0%r ]
M3.noisy_sum ~ M3.noisy_sum
  : adjacent_e ls{1} ls{2} j /\ n=size ls{1}
    ==> res{2} = res{1} ].
proof.
move => H. proc.
seq 3 3: (adjacent_e ls{1} ls{2} j /\ = {i, s, output} /\ i{1} = 0
        /\ s{1} = 0 /\ n=size ls{1}).
toequiv; auto.
(* notice the budget function we are using for epsilon *)
awhile [ (fun x => if j=n-x -1 then eps else 0%r ) & (fun _ => 0%r) ] n [n -i-1]
(adjacent_e ls{1} ls{2} j /\ = {i, output} /\ 0 <= i{1} <= size ls{1} /\ n=size ls{1} /\
 (i{1}=j =>
`|(nth 0 ls{1} i{1}) - (nth 0 ls{2} i{2})| <= 1 /\
eq_in_range ls{1} ls{2} (i{1}+1) (size ls{1} - 1)) /\
0 <= size ls{1}); first 4 try (auto; progress; smt(ge0_eps)).
search pred1. search iota_. rewrite /bigi.
rewrite -big_mkcond. apply bigi_eps => //. rewrite sumr_const intmulr; smt().
move => k.
have H1: (j = n - k - 1) \/\ (j < n - k - 1). smt().
elim H1 => [?!?].
wp. progress.
lap 0 1. smt().
progress. smt().
progress. smt().
smt(). smt().
smt(adjacent_sub_abs_bound).
smt(adjacent_ne_sub_eq_after).
smt(size_eq_adjacent).
wp.
lap 0 0. smt().
progress.
smt(size_eq_adjacent).
auto; progress. smt().
smt().
smt(adjacent_sub_abs_bound).
smt(adjacent_ne_sub_eq_after).
smt(size_eq_adjacent).
smt(size_eq_adjacent).
qed.
```

Pointwise Differential Privacy

Definition

Given $\epsilon, \delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ϵ, δ) -differentially private iff

for all adjacent database D, D' and for every $S \subseteq R$:

$$\Pr[Q(D) \in S] \leq \exp(\epsilon) \Pr[Q(D') \in S] + \delta$$

Pointwise Definition

Given $\epsilon, \delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ϵ, δ) -differentially private iff

for all adjacent database D, D' and for every $r \in R$, we have δ_r such that:

$$\Pr[Q(D) \in S] \leq \exp(\epsilon) \Pr[Q(D') \in S] + \delta_r$$

and $\sum \delta_r \leq \delta$

Above Threshold

```
aboveT (db :int list,n:int,t:int):int = {
  s<-0;
  i <- 0;
  r <- -1;
  nT <$ lap (eps/4%r) t;
  while (i < n) {
    s <$ lap (eps/2%r) (evalQ i db);
    if (nT < s /\ r = -1){
      r <- i;
    }
    i <- i + 1;
  }
  return r;
}
```

Pointwise DP in Aprhl

forall $r \in R$

$\vdash_{\varepsilon, \delta r} C_1 \sim C_2 : P \implies x \langle 1 \rangle = r \implies x \langle 2 \rangle = r$

$$\sum \delta r \leq \delta$$

$\vdash_{\varepsilon, \delta} C_1 \sim C_2 : P \implies x \langle 1 \rangle = x \langle 2 \rangle$

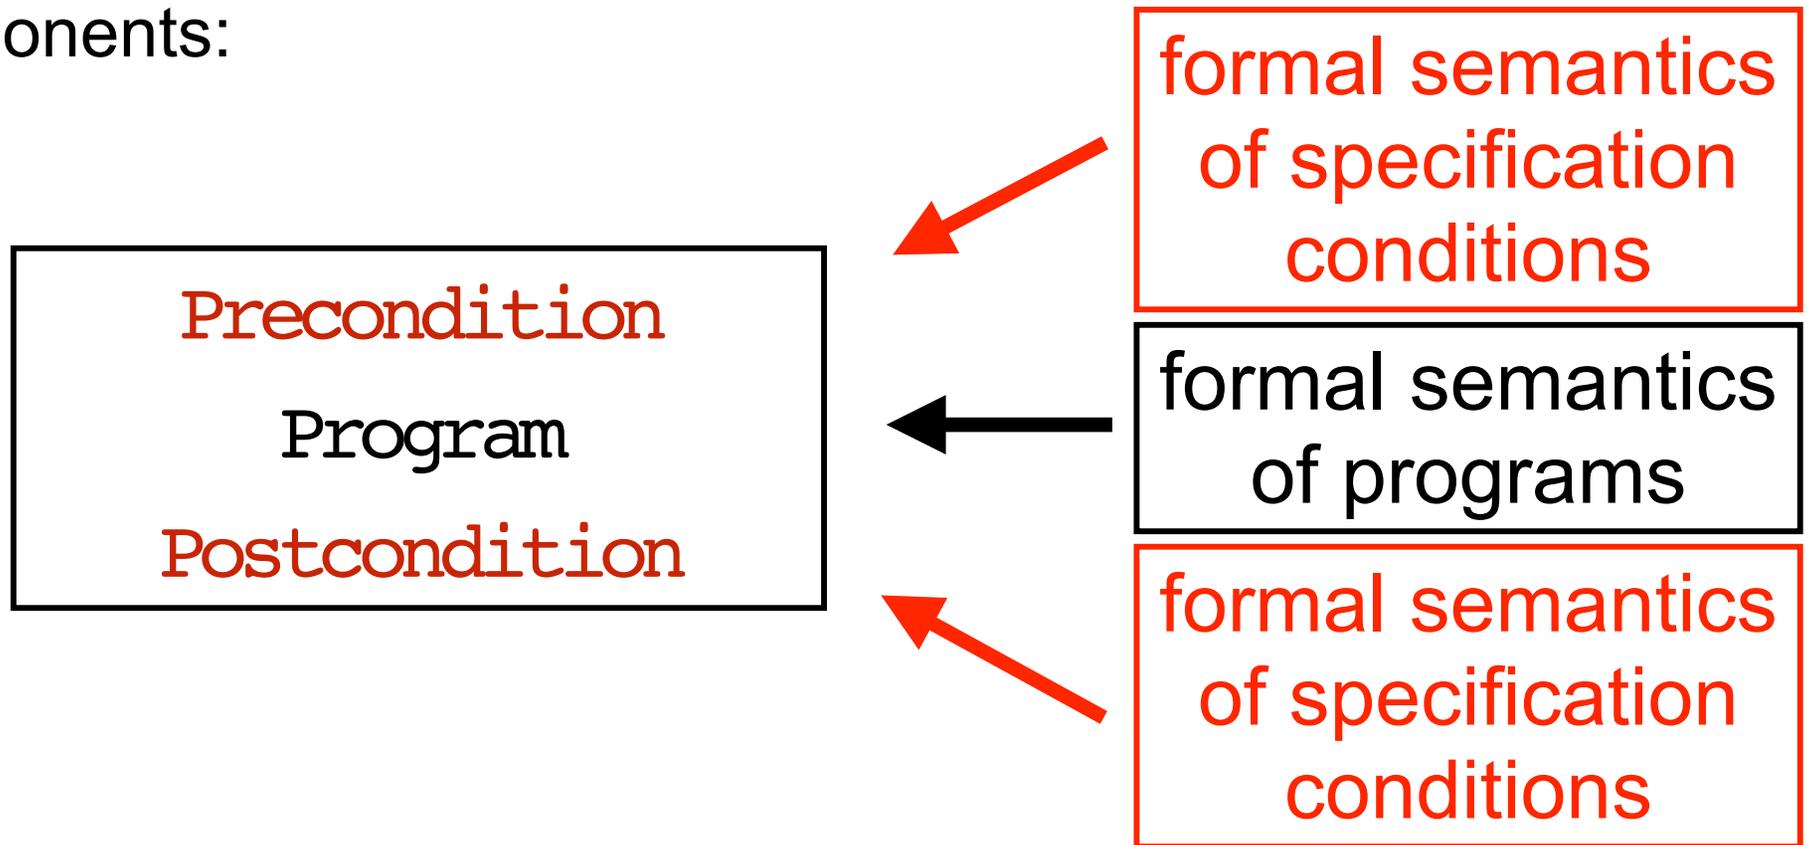
The corresponding tactic in EC is

$\text{pweq}(r, r)$

Review of the class

Formal Semantics

We need to assign a formal meaning to the different components:



We also need to describe the rules which combine program and specifications.

Programming Language

```
c ::= abort
    | skip
    | x := e
    | c ; c
    | if e then c else c
    | while e do c
```

x, y, z, \dots program variables

e_1, e_2, \dots expressions

c_1, c_2, \dots commands

Summary of the Semantics of Commands

$$\{\text{abort}\}_m = \perp$$

$$\{\text{skip}\}_m = m$$

$$\{x := e\}_m = m[x \leftarrow \{e\}_m]$$

$$\{c; c'\}_m = \{c'\}_{m'} \quad \text{If } \{c\}_m = m'$$

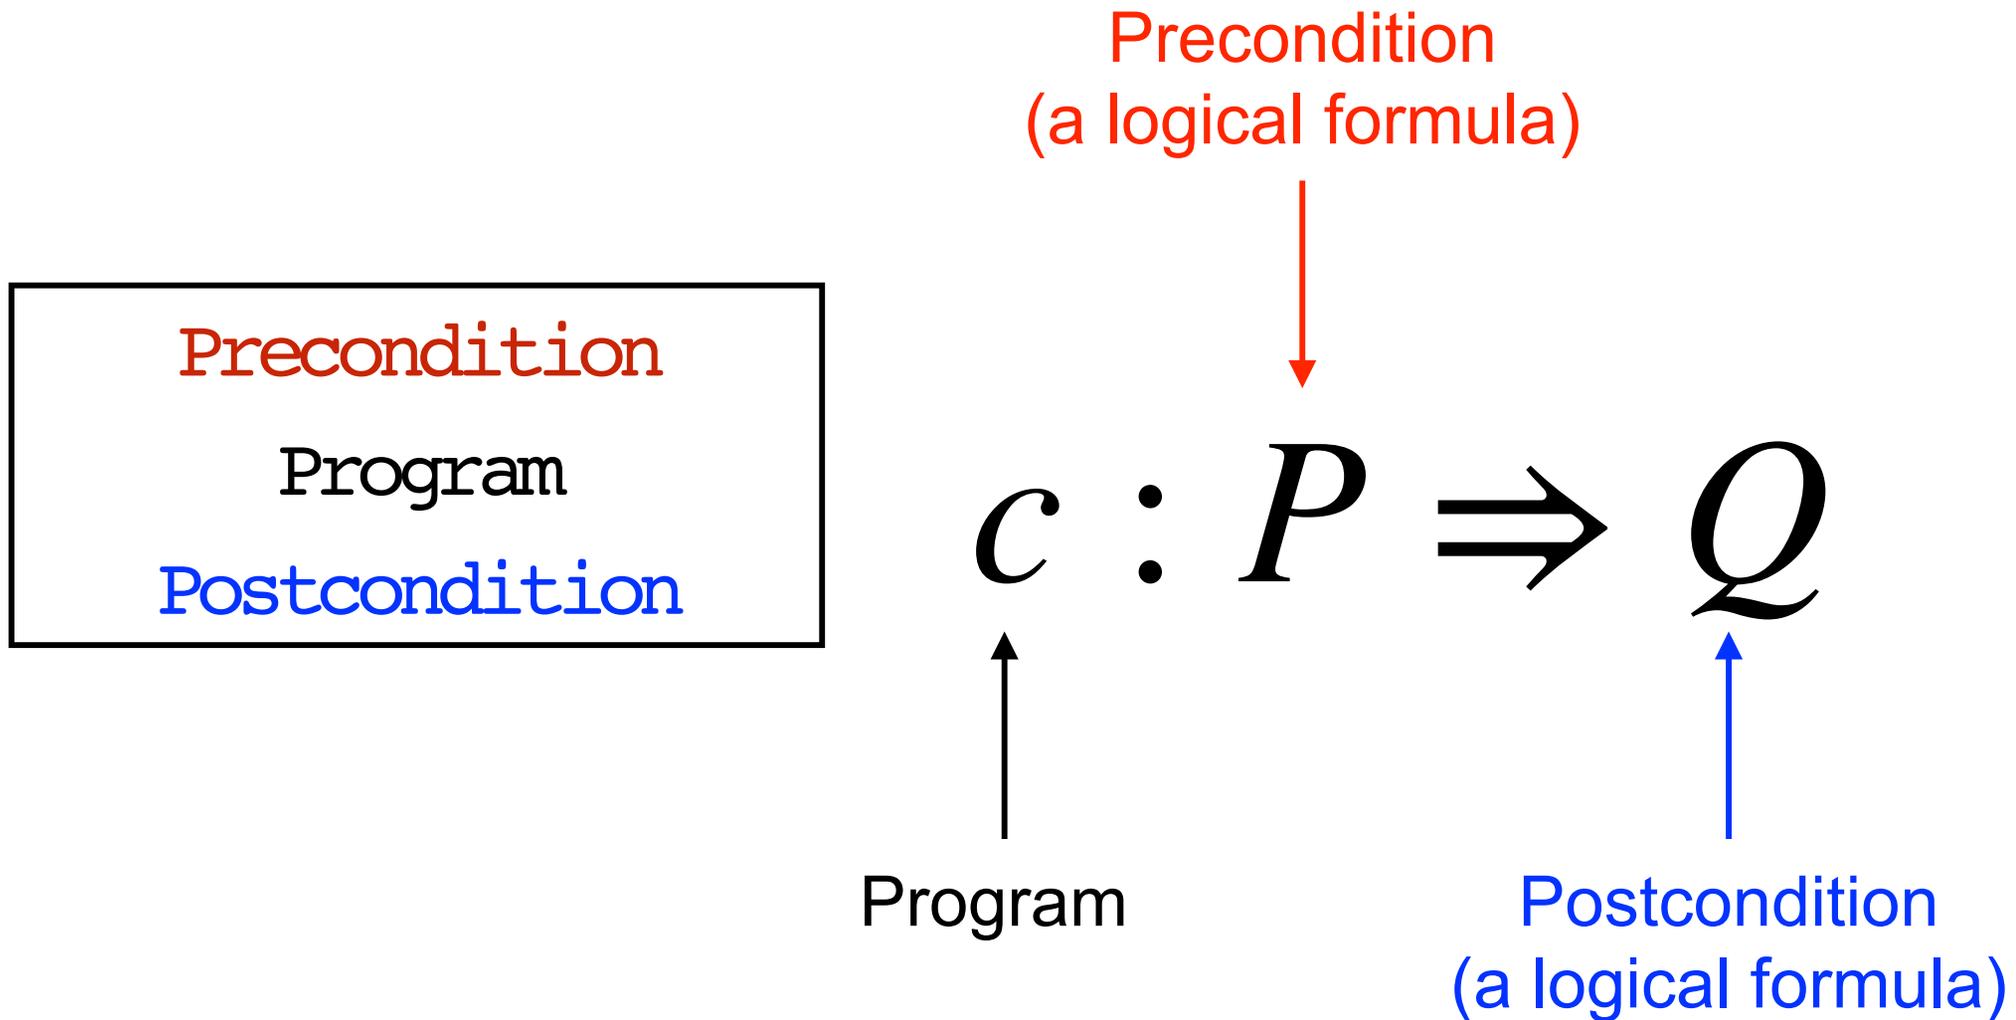
$$\{c; c'\}_m = \perp \quad \text{If } \{c\}_m = \perp$$

$$\{\text{if } e \text{ then } c_t \text{ else } c_f\}_m = \{c_t\}_m \quad \text{If } \{e\}_m = \text{true}$$

$$\{\text{if } e \text{ then } c_t \text{ else } c_f\}_m = \{c_f\}_m \quad \text{If } \{e\}_m = \text{false}$$

$$\{\text{while } e \text{ do } c\}_m = \sup_{n \in \text{Nat}} \{\text{while}_n e \text{ do } c\}_m$$

Hoare triple



Validity of Hoare triple

We say that the triple $c : P \Rightarrow Q$ is **valid**
if and only if

for every memory m such that $P(m)$
and memory m' such that $\{c\}_m = m'$
we have $Q(m')$.

Is this condition easy to check?

An example

\vdash

<pre>x := 3; y := 1; while x > 1 do y := y + 1; x := x - 1;</pre>
--

$: \{true\} \Rightarrow \{y = 3\}$

Soundness

If we can derive $\vdash c : P \Rightarrow Q$ through the rules of the logic, then the triple $c : P \Rightarrow Q$ is valid.

Relative Completeness

$$P \Rightarrow S \quad \vdash c : S \Rightarrow R \quad R \Rightarrow Q$$

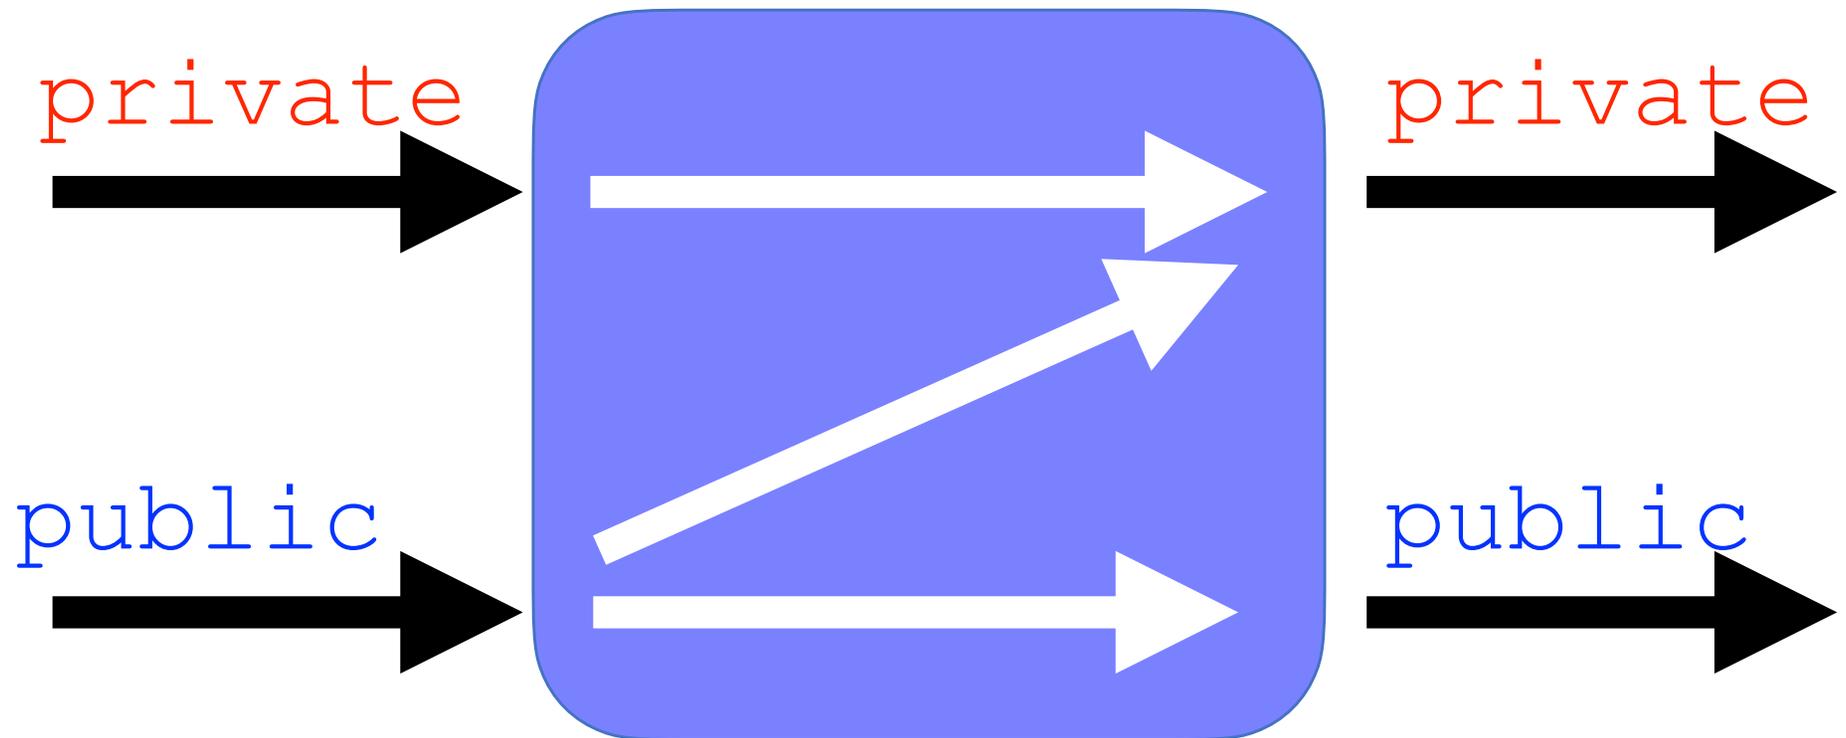
$$\vdash c : P \Rightarrow Q$$

If a triple $c : P \Rightarrow Q$ is valid, and we have an oracle to derive all the true statements of the form $P \Rightarrow S$ and of the form $R \Rightarrow Q$, then we can derive $\vdash c : P \Rightarrow Q$ through the rules of the logic.

Noninterference

In symbols

$m_1 \sim_{\text{low}} m_2$ and $\{c\}_{m_1} = m_1'$ and $m_2' \{c\}_{m_2} = m_2'$
implies $m_1' \sim_{\text{low}} m_2'$



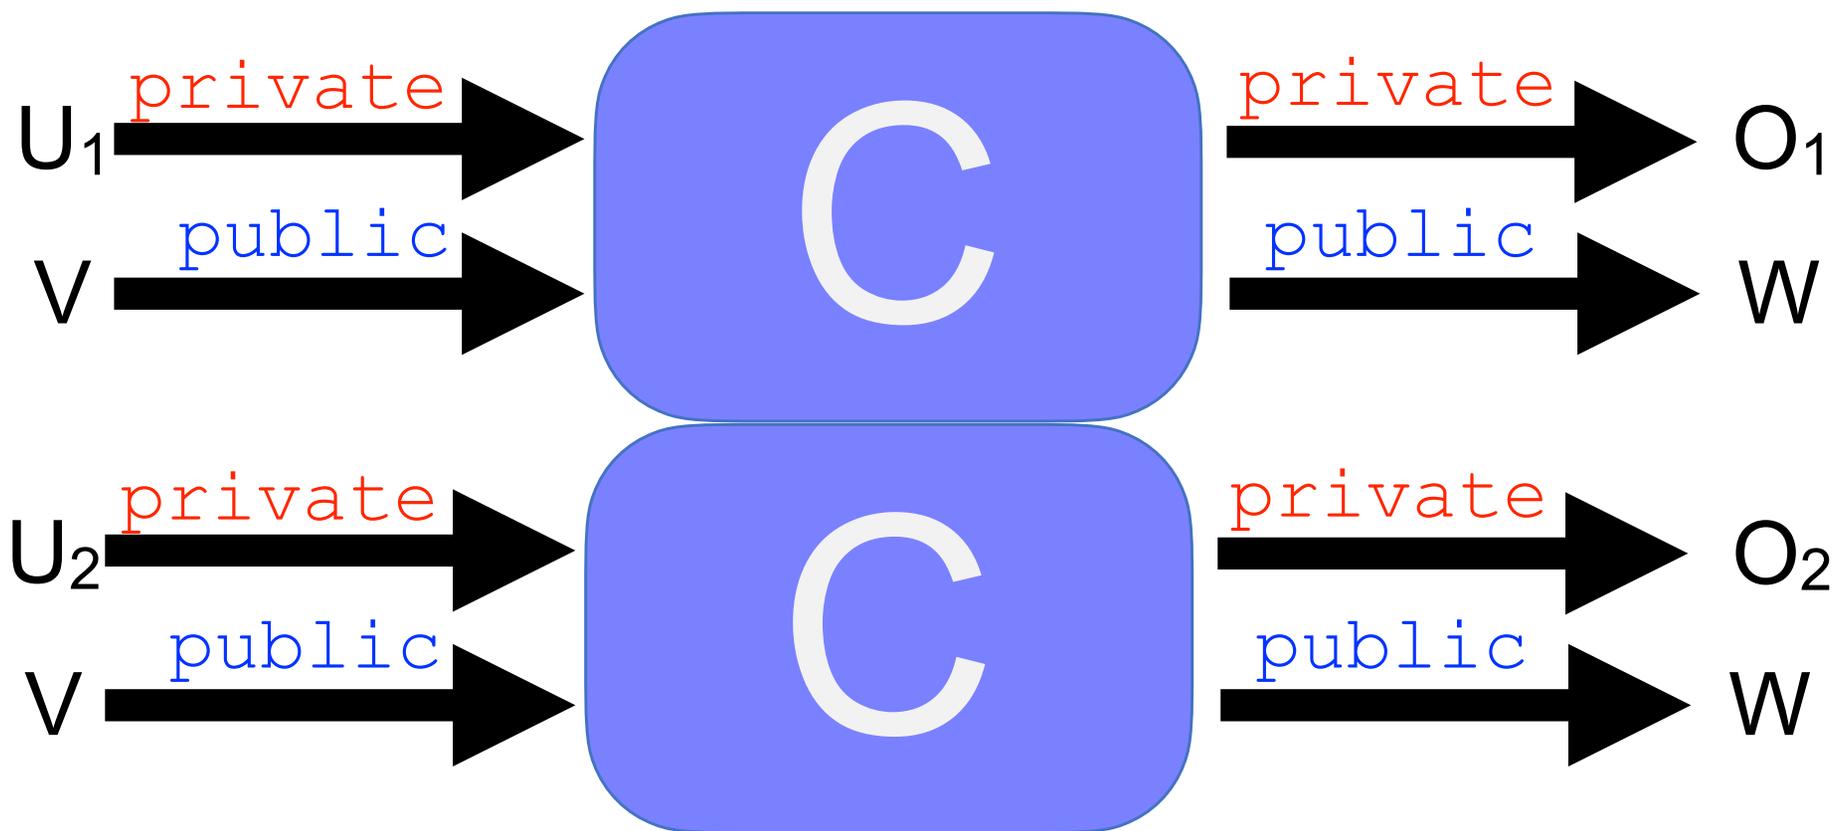
Relational Property

In symbols, c is **noninterferent** if and only if

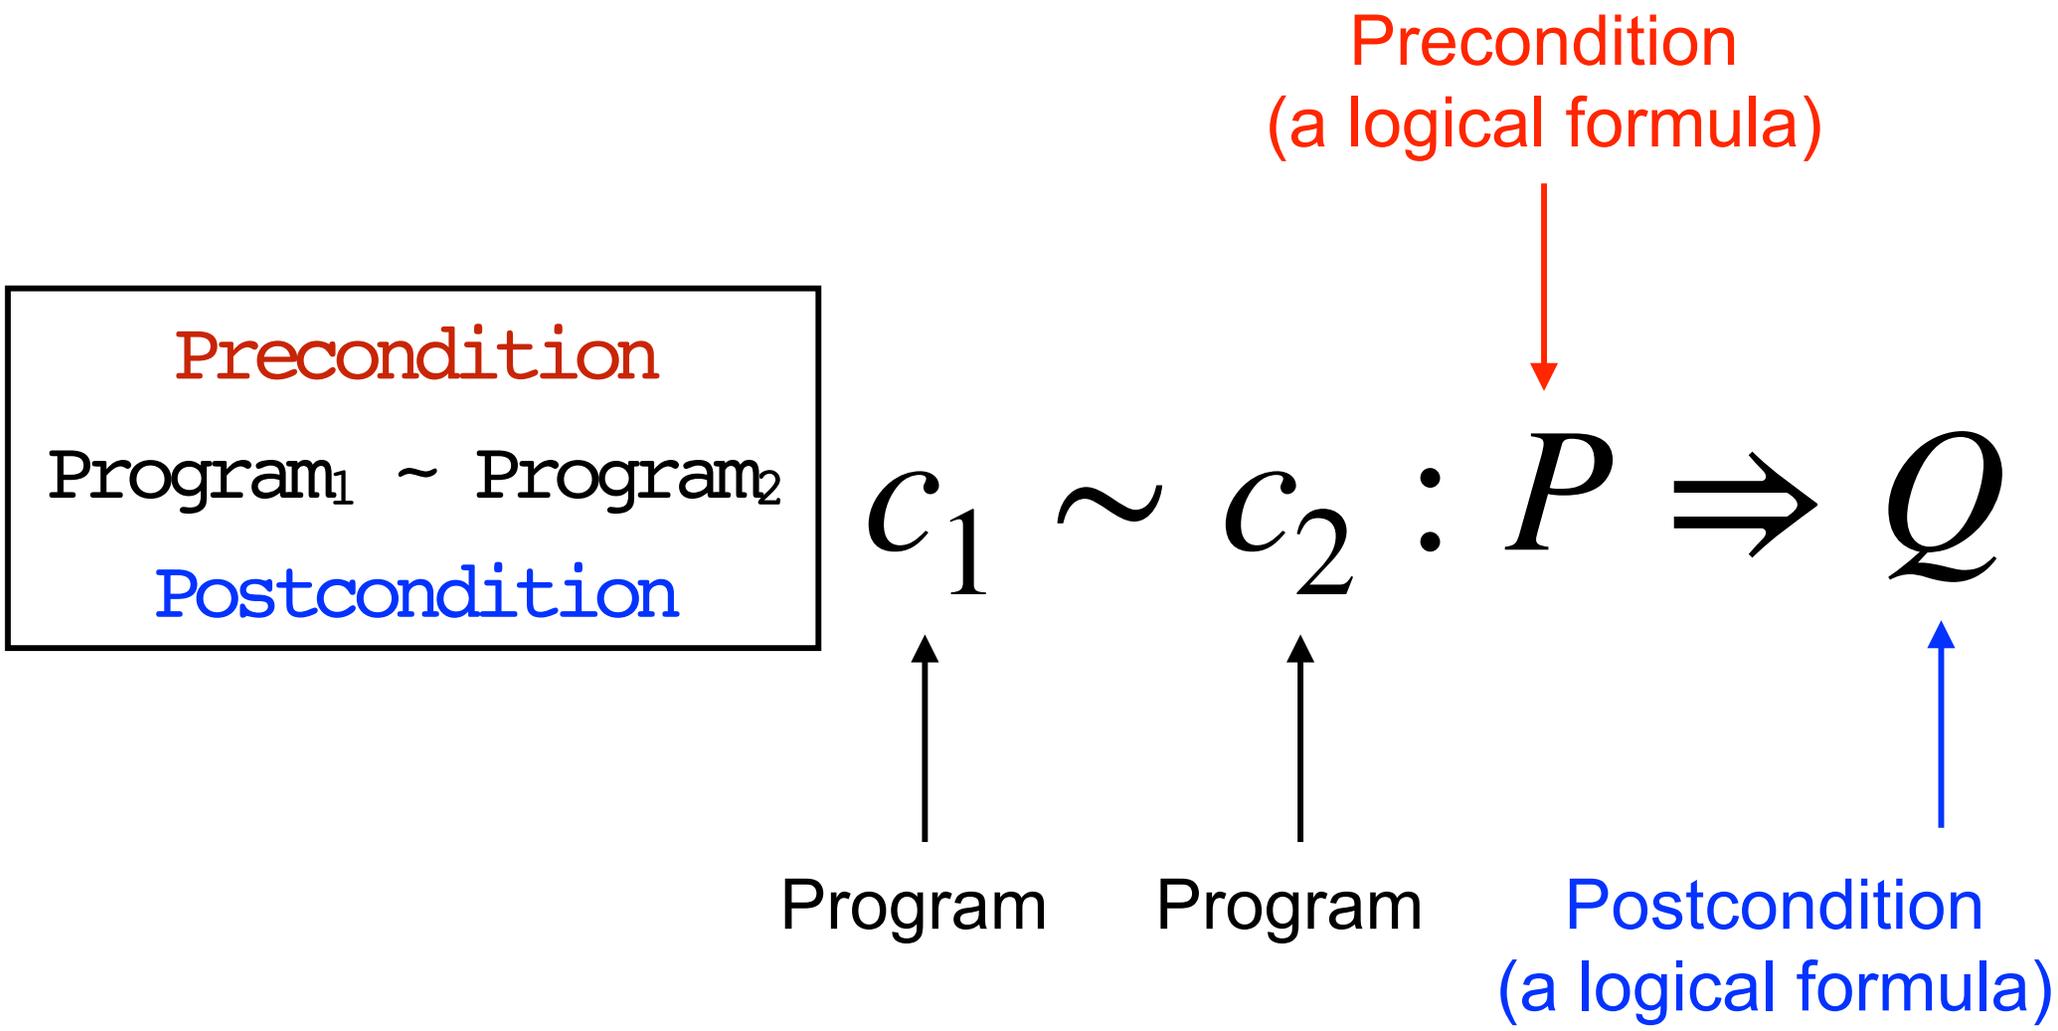
for every $m_1 \sim_{\text{low}} m_2$:

1) $\{c\}_{m_1} = \perp$ iff $\{c\}_{m_2} = \perp$

2) $\{c\}_{m_1} = m_1'$ and $\{c\}_{m_2} = m_2'$ implies $m_1' \sim_{\text{low}} m_2'$



Relational Hoare Logic - RHL



Validity of Hoare quadruple

We say that the quadruple $c_1 \sim c_2 : P \Rightarrow Q$ is **valid** if and only if for every pair of memories m_1, m_2 such that $P(m_1, m_2)$ we have:

- 1) $\{c_1\}_{m_1} = \perp$ iff $\{c_2\}_{m_2} = \perp$
- 2) $\{c_1\}_{m_1} = m_1'$ and $\{c_2\}_{m_2} = m_2'$ implies $Q(m_1', m_2')$.

Is this easy to check?

An example

```
s1:public
s2:private
r:private
i:public

proc Compare (s1:list[n] bool,s2:list[n] bool)
i:=0;
r:=0;
while i<n do
  if not(s1[i]=s2[i]) then
    r:=1
  i:=i+1

: n>0 /\ =low ⇒ =low
```

Soundness and completeness with respect to Hoare Logic

$$\vdash_{\text{RHL}} C_1 \sim C_2 : P \Rightarrow Q$$

iff

$$\vdash_{\text{HL}} C_1 ; C_2 : P \Rightarrow Q$$

Under the assumption that we can partition the memory adequately, and that we have termination.

An example

```
OneTimePad(m : private msg) : public msg  
  key := $ Uniform({0,1}n);  
  cipher := msg xor key;  
  return cipher
```

Learning a ciphertext does not change any a priori knowledge about the likelihood of messages.

Probabilistic While (PWhile)

```
c ::= abort
    | skip
    | x := e
    | x :=$ d
    | c ; c
    | if e then c else c
    | while e do c
```

d_1, d_2, \dots probabilistic expressions

Semantics of Commands

This is defined on the structure of commands:

$$\{\text{abort}\}_m = \mathbf{0}$$

$$\{\text{skip}\}_m = \text{unit}(m)$$

$$\{x := e\}_m = \text{unit}(m[x \leftarrow \{e\}_m])$$

$$\{x := \$ d\}_m = \text{let } a = \{d\}_m \text{ in } \text{unit}(m[x \leftarrow a])$$

$$\{c; c'\}_m = \text{let } m' = \{c\}_m \text{ in } \{c'\}_{m'}$$

$$\{\text{if } e \text{ then } c_t \text{ else } c_f\}_m = \{c_t\}_m \text{ if } \{e\}_m = \text{true}$$

$$\{\text{if } e \text{ then } c_t \text{ else } c_f\}_m = \{c_f\}_m \text{ if } \{e\}_m = \text{false}$$

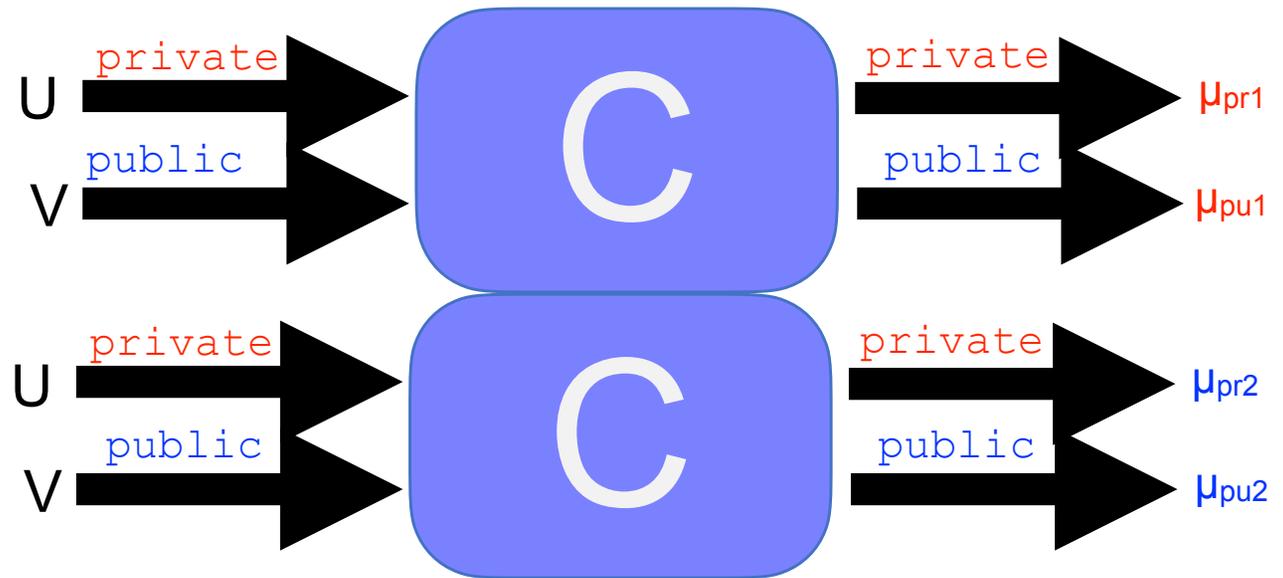
$$\{\text{while } e \text{ do } c\}_m = \sup_{n \in \text{Nat}} \mu_n$$

$$\mu_n =$$

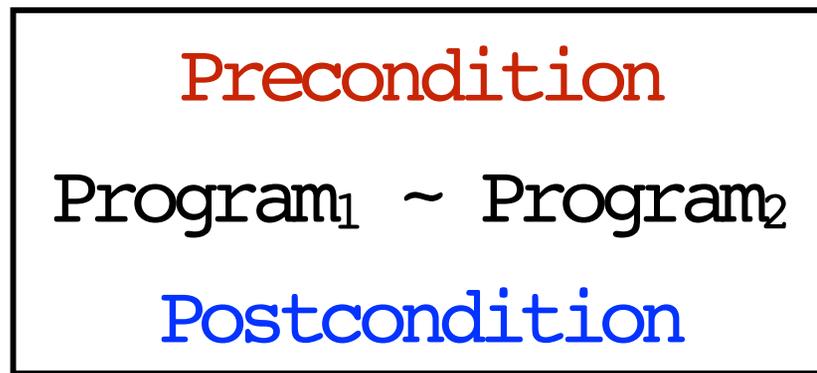
$$\text{let } m' = \{(\text{while}^n e \text{ do } c)\}_m \text{ in } \{\text{if } e \text{ then abort}\}_{m'}$$

Probabilistic Noninterference as a Relational Property

c is **probabilistically noninterferent** if and only if for every $m_1 \sim_{\text{low}} m_2$:
 $\{c\}_{m_1} = \mu_1$ and $\{c\}_{m_2} = \mu_2$ implies $\mu_1 \sim_{\text{low}} \mu_2$



Probabilistic Relational Hoare Quadruples



Precondition
(a logical formula)



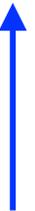
$$c_1 \sim c_2 : P \Rightarrow Q$$



Probabilistic
Program



Probabilistic
Program



Postcondition
(a logical formula)

Validity of Probabilistic Hoare quadruple

We say that the quadruple $c_1 \sim c_2 : P \Rightarrow Q$ is **valid** if and only if for every pair of memories m_1, m_2 such that $P(m_1, m_2)$ we have:

$\{c_1\}_{m_1} = \mu_1$ and $\{c_2\}_{m_2} = \mu_2$ implies

$Q^*(\mu_1, \mu_2)$.

Relational lifting of a predicate

We say that two subdistributions $\mu_1 \subseteq D(A)$ and $\mu_2 \subseteq D(B)$ are in the relational lifting of the relation $R \subseteq A \times B$, denoted $\mu_1 R^* \mu_2$ if and only if there exist a subdistribution $\mu \subseteq D(A \times B)$ such that:

- 1) if $\mu(a, b) > 0$, then $(a, b) \in R$.
- 2) $\pi_1(\mu) = \mu_1$ and $\pi_2(\mu) = \mu_2$

R-Coupling

Given two distributions $\mu_1 \in D(A)$, and $\mu_2 \in D(B)$, an **R-coupling** between them, for $R \subseteq A \times B$, is a joint distribution $\mu \in D(A \times B)$ such that:

- 1) the marginal distributions of μ are μ_1 and μ_2 , respectively,
- 2) the support of μ is contained in R . That is, if $\mu(a, b) > 0$, then $(a, b) \in R$.

A sufficient condition for R-Coupling

Given two distributions $\mu_1 \in \mathcal{D}(A)$, and $\mu_2 \in \mathcal{D}(B)$, and a relation $R \subseteq A \times B$, if there is a mapping $h: A \rightarrow B$ such that:

- 1) h is a bijective map between elements in $\text{supp}(\mu_1)$ and $\text{supp}(\mu_2)$,
- 2) for every $a \in \text{supp}(\mu_1)$, $(a, h(a)) \in R$
- 3) $\Pr_{x \sim \mu_1} [x = a] = \Pr_{x \sim \mu_2} [x = h(a)]$

Then, there is an **R-coupling** between μ_1 and μ_2 .
We write $h \triangleleft (\mu_1, \mu_2)$ in this case.

Probabilistic Relational Hoare Logic

Random Assignment

$$h \triangleleft (\{ d_1 \} , \{ d_2 \})$$
$$P = \forall v, v \in \text{supp} (\{ d_1 \})$$
$$\Rightarrow Q [v / x_1 \langle 1 \rangle , h (v) / x_2 \langle 2 \rangle]$$

$$\vdash x_1 := \$ d_1 \sim x_2 := \$ d_2 : P \Rightarrow Q$$

Consequences of Coupling

Given the following pRHL judgment

$$\vdash c_1 \sim c_2 : \text{True} \Rightarrow Q$$

We have that:

if $Q \Rightarrow (R\langle 1 \rangle \iff S\langle 2 \rangle)$, then $\Pr[c_1 : R] = \Pr[c_2 : S]$

if $Q \Rightarrow (R\langle 1 \rangle \Rightarrow S\langle 2 \rangle)$, then $\Pr[c_1 : R] \leq \Pr[c_2 : S]$

A more realistic example

```
StreamCipher (m : private msg[n]) : public msg[n]  
  pkey := $ PRG (Uniform({0,1}k));  
  cipher := msg xor pkey;  
  return cipher
```

Properties of PRG

We would like the PRG to increase the number of random bits but also to guarantee the result to be (almost) random.

We can express this as:

$$\text{PRG: } \{0,1\}^k \rightarrow \{0,1\}^n \text{ for } n > k$$

$$\Delta(\text{PRG}(\text{Uniform}(\{0,1\}^k), \text{Uniform}(\{0,1\}^n)) \leq 2^{-n}$$

In fact this is a too strong requirement - usually we require that every polynomial time adversary cannot distinguish the two distributions in statistical distance

Approximate Probabilistic Relational Hoare Logic

Indistinguishability
parameter

Precondition
(a logical formula)

$$\vdash \delta \ C_1 \sim C_2 : P \Rightarrow Q$$

The diagram shows the formula $\vdash \delta \ C_1 \sim C_2 : P \Rightarrow Q$ with five arrows pointing to its components: a purple arrow points to δ , a black arrow points to C_1 , a black arrow points to C_2 , a red arrow points to P , and a blue arrow points to Q .

Probabilistic
Program

Probabilistic
Program

Postcondition
(a logical formula)

R- δ -Coupling

Given two distributions $\mu_1 \in D(A)$, and $\mu_2 \in D(B)$, we have an R- δ -coupling between them, for $R \subseteq A \times B$ and $0 \leq \delta \leq 1$, if there are two joint distributions $\mu_L, \mu_R \in D(A \times B)$ such that:

- 1) $\pi_1(\mu_L) = \mu_1$ and $\pi_2(\mu_R) = \mu_2$,
- 2) the support of μ_L and μ_R is contained in R .
That is, if $\mu_L(a, b) > 0$, then $(a, b) \in R$,
and if $\mu_R(a, b) > 0$, then $(a, b) \in R$.
- 3) $\Delta(\mu_L, \mu_R) \leq \delta$

Example of R- δ -Coupling

μ_1

00	0.2
01	0.25
10	0.25
11	0.3

μ_2

00	0
01	0.40
10	0
11	0.6

$$R(a, b) = \{ a \leq b \}$$

μ_L	00	01	10	11
00		0.20		
01		0.25		
10				0.25
11				0.30

μ_R	00	01	10	11
00		0.20		
01		0.20		
10				0.3
11				0.3

$$\Delta(\mu_L, \mu_R) = 0.05$$

Approximate relational lifting of a predicate

We say that two subdistributions $\mu_1 \subseteq D(A)$ and $\mu_2 \subseteq D(B)$ are in the relational δ -lifting of the relation $R \subseteq A \times B$, denoted $\mu_1 R_\delta^* \mu_2$ if and only if there exist an R -coupling between them.

Validity of approximate Probabilistic Hoare judgments

We say that the quadruple $\vdash_{\delta} c_1 \sim c_2 : P \Rightarrow Q$ is **valid** if and only if for every pair of memories m_1, m_2 such that $P(m_1, m_2)$ we have:

$\{c_1\}_{m_1} = \mu_1$ and $\{c_2\}_{m_2} = \mu_2$ implies

$Q_{\delta^*}(\mu_1, \mu_2)$.

Releasing the mean of Some Data

```
Mean (d : private data) : public real
  i:=0;
  s:=0;
  while (i<size(d))
    s:=s + d[i]
    i:=i+1;
  return (s/i)
```

(ϵ, δ) -Differential Privacy

Definition

Given $\epsilon, \delta \geq 0$, a probabilistic query $Q: X^n \rightarrow R$ is (ϵ, δ) -differentially private iff for all adjacent databases b_1, b_2 and for every $S \subseteq R$:

$$\Pr[Q(b_1) \in S] \leq \exp(\epsilon) \Pr[Q(b_2) \in S] + \delta$$

(ε, δ) -indistinguishability

We can define a ε -skewed version of statistical distance. We call this notion ε -distance.

$$\Delta_\varepsilon(\mu_1, \mu_2) = \sup_{E \subseteq A} \max(\mu_1(E) - e^\varepsilon \mu_2(E), \mu_2(E) - e^\varepsilon \mu_1(E), 0)$$

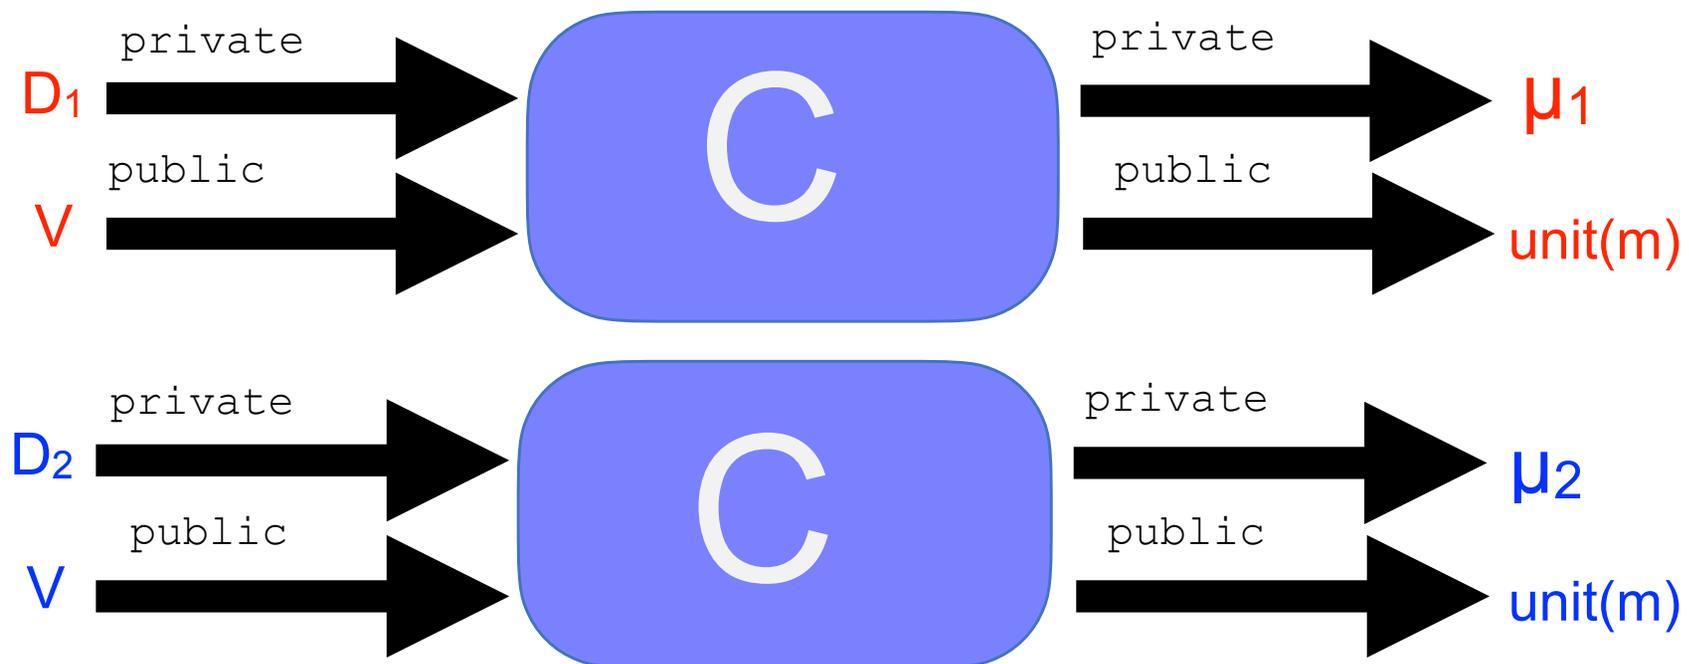
We say that two distributions $\mu_1, \mu_2 \in D(A)$, are at (ε, δ) -indistinguishable if:

$$\Delta_\varepsilon(\mu_1, \mu_2) \leq \delta$$

Differential Privacy as a Relational Property

c is **differentially private** if and only if for every $m_1 \sim m_2$ (extending the notion of adjacency to memories):

$\{c\}_{m_1} = \mu_1$ and $\{c\}_{m_2} = \mu_2$ implies $\Delta_\epsilon(\mu_1, \mu_2) \leq \delta$



apRHL

Indistinguishability
parameter

Precondition
(a logical formula)

$$\vdash_{\epsilon, \delta} C_1 \sim C_2 : P \Rightarrow Q$$

Probabilistic
Program

Probabilistic
Program

Postcondition
(a logical formula)

Validity of apRHL judgments

We say that the quadruple $\vdash_{\varepsilon, \delta} c_1 \sim c_2 : P \Rightarrow Q$ is **valid** if and only if for every pair of memories m_1, m_2 such that $P(m_1, m_2)$ we have:

$\{c_1\}_{m_1} = \mu_1$ and $\{c_2\}_{m_2} = \mu_2$ implies $Q_{\varepsilon, \delta}^*(\mu_1, \mu_2)$.

R- ε - δ -Coupling

Given two distributions $\mu_1 \in \mathcal{D}(A)$, and $\mu_2 \in \mathcal{D}(B)$, we have an R- ε - δ -coupling between them, for $R \subseteq A \times B$ and $0 \leq \varepsilon$ and $0 \leq \delta \leq 1$, if there are two joint distributions $\mu_L, \mu_R \in \mathcal{D}(A \times B)$ such that:

- 1) $\pi_1(\mu_L) = \mu_1$ and $\pi_2(\mu_R) = \mu_2$,
- 2) the support of μ_L and μ_R is contained in R .
That is, if $\mu_L(a, b) > 0$, then $(a, b) \in R$,
and if $\mu_R(a, b) > 0$, then $(a, b) \in R$.
- 3) $\Delta_\varepsilon(\mu_L, \mu_R) \leq \delta$

Symmetric Encryption Schemes

- Our treatment of symmetric encryption schemes is parameterized by three types:

```
type key. (* encryption keys, key_len bits *)
```

```
type text. (* plaintexts, text_len bits *)
```

```
type cipher. (* ciphertexts - scheme specific *)
```

- An encryption scheme is a *stateless* implementation of this module interface:

```
module type ENC = {
```

```
  proc key_gen() : key (* key generation *)
```

```
  proc enc(k : key, x : text) : cipher (* encryption *)
```

```
  proc dec(k : key, c : cipher) : text (* decryption *)
```

```
};
```

IND-CPA Game

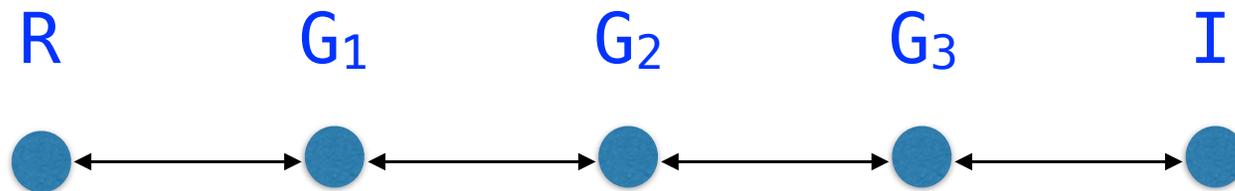
- The IND-CPA Game is parameterized by an encryption scheme and an encryption adversary:

```
module INDCPA (Enc : ENC, Adv : ADV) = {  
  module E0 = Enc0(Enc)          (* make E0 from Enc *)  
  module A = Adv(E0)            (* connect Adv to E0 *)  
  proc main() : bool = {  
    var b, b' : bool; var x1, x2 : text; var c : cipher;  
    E0.init();                   (* initialize E0 *)  
    (x1, x2) <@ A.choose();      (* let A choose x1/x2 *)  
    b <$ {0,1};                  (* choose boolean b *)  
    c <@ E0.genc(b ? x1 : x2);   (* encrypt x1 or x2 *)  
    b' <@ A.guess(c);           (* let A guess b from c *)  
    return b = b';              (* see if A won *)  
  }  
}.
```

Sequence of Games Approach

- Our proof of IND-CPA security uses the *sequence of games approach*, which is used to connect a “real” game **R** with an “ideal” game **I** via a sequence of intermediate games.
- Each of these games is parameterized by the adversary, and each game has a **main** procedure returning a boolean.
- We want to establish an upper bound for

$$\left| \Pr[\mathbf{R}.\text{main}() @ \&m : \text{res}] - \Pr[\mathbf{I}.\text{main}() : \text{res}] \right|$$



Sequence of Games Approach

- Suppose we can prove

$$\text{` | Pr}[R.main() @ \&m : res] - \text{Pr}[G_1.main() : res] | \leq b_1$$

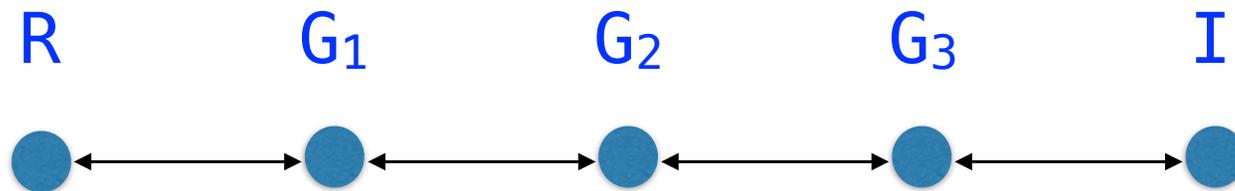
$$\text{` | Pr}[G_1.main() @ \&m : res] - \text{Pr}[G_2.main() : res] | \leq b_2$$

$$\text{` | Pr}[G_2.main() @ \&m : res] - \text{Pr}[G_3.main() : res] | \leq b_3$$

$$\text{` | Pr}[G_3.main() @ \&m : res] - \text{Pr}[I.main() : res] | \leq b_4$$

for some b_1 , b_2 , b_3 and b_4 . Then we can conclude

$$\text{` | Pr}[R.main() @ \&m : res] - \text{Pr}[I.main() @ \&m : res] | \leq b_1 + b_2 + b_3 + b_4$$



Step 1: Replacing PRF with TRF

- In our first step, we switch to using a true random function instead of a pseudorandom function in our encryption scheme.
 - We have an exact model of how the TRF works.
- When doing this, we inline the encryption scheme into a new kind of encryption oracle, **E0_RF**, which is parameterized by a random function.
- We also instrument **E0_RF** to detect two kinds of “clashes” (repetitions) in the generation of the inputs to the random function.
 - This is in preparation for Steps 2 and 3.

Step 2: Oblivious Update in **genc**

- In Step 2, we make use of **up to bad reasoning**, to transition to a game in which the encryption oracle, **E0_0**, uses a true random function and “obliviously” (“O” for “oblivious”) updates the true random function’s map — i.e., overwrites what may already be stored in the map.

Step 3: Independent Choice in `genc`

- In Step 3, we again make use of up to bad reasoning, this time transitioning to a game in which the encryption oracle, `E0_I`, chooses the text value to be exclusive or-ed with the plaintext in a way that is “independent” (“I” for “independent”) from the true random function’s map, i.e., without updating that map.
- We no longer need to detect “pre” clashes (clashes in `genc` with a `u` chosen in a call to `enc_pre`).

Step 4: One-time Pad Argument

- In Step 4, we can switch to an encryption oracle `E0_N` in which the right side of the ciphertext produced by `E0_N.genc` makes no (“N” for “no”) reference to the plaintext.
- We no longer need any instrumentation for detecting clashes.

Step 5: Proving G4's Probability

- When proving

```
local lemma G4_prob &m :
```

```
  Pr[G4.main() @ &m : res] = 1%r / 2%r.
```

we can reorder

```
b <$ {0,1};  
c <@ E0_N.genc(text0);  
b' <@ A.guess(c);  
return b = b';
```

to

```
c <@ E0_N.genc(text0);  
b' <@ A.guess(c);  
b <$ {0,1};  
return b = b';
```

- We use that Adv's procedures are lossless.

IND-CPA Security Result

```
lemma INDCPA (Adv <: ADV{Enc0, PRF, TRF, Adv2RFA}) &m :
  (forall (E0 <: E0{Adv}),
    islossless E0.enc_pre => islossless Adv(E0).choose) =>
  (forall (E0 <: E0{Adv}),
    islossless E0.enc_post => islossless Adv(E0).guess) =>
  `|Pr[INDCPA(Enc, Adv).main() @ &m : res] -
    1%r / 2%r| <=
  `|Pr[GRF(PRF, Adv2RFA(Adv)).main() @ &m : res] -
    Pr[GRF(TRF, Adv2RFA(Adv)).main() @ &m : res]| +
  (limit_pre%r + limit_post%r) / (2 ^ text_len)%r.
```

- Q: If we remove the restriction on **Adv** (**{Enc0, PRF, TRF, Adv2RFA}**), what would happen?
- A: Various tactic applications would fail; e.g., calls to the **Adv**'s procedures, as they could invalidate assumptions.

Any Question?

