

CMSC/Math 456: Cryptography (Fall 2022)

Lecture 14

Daniel Gottesman

Administrative

Midterm: Thursday, Oct. 19 (1 week from today)

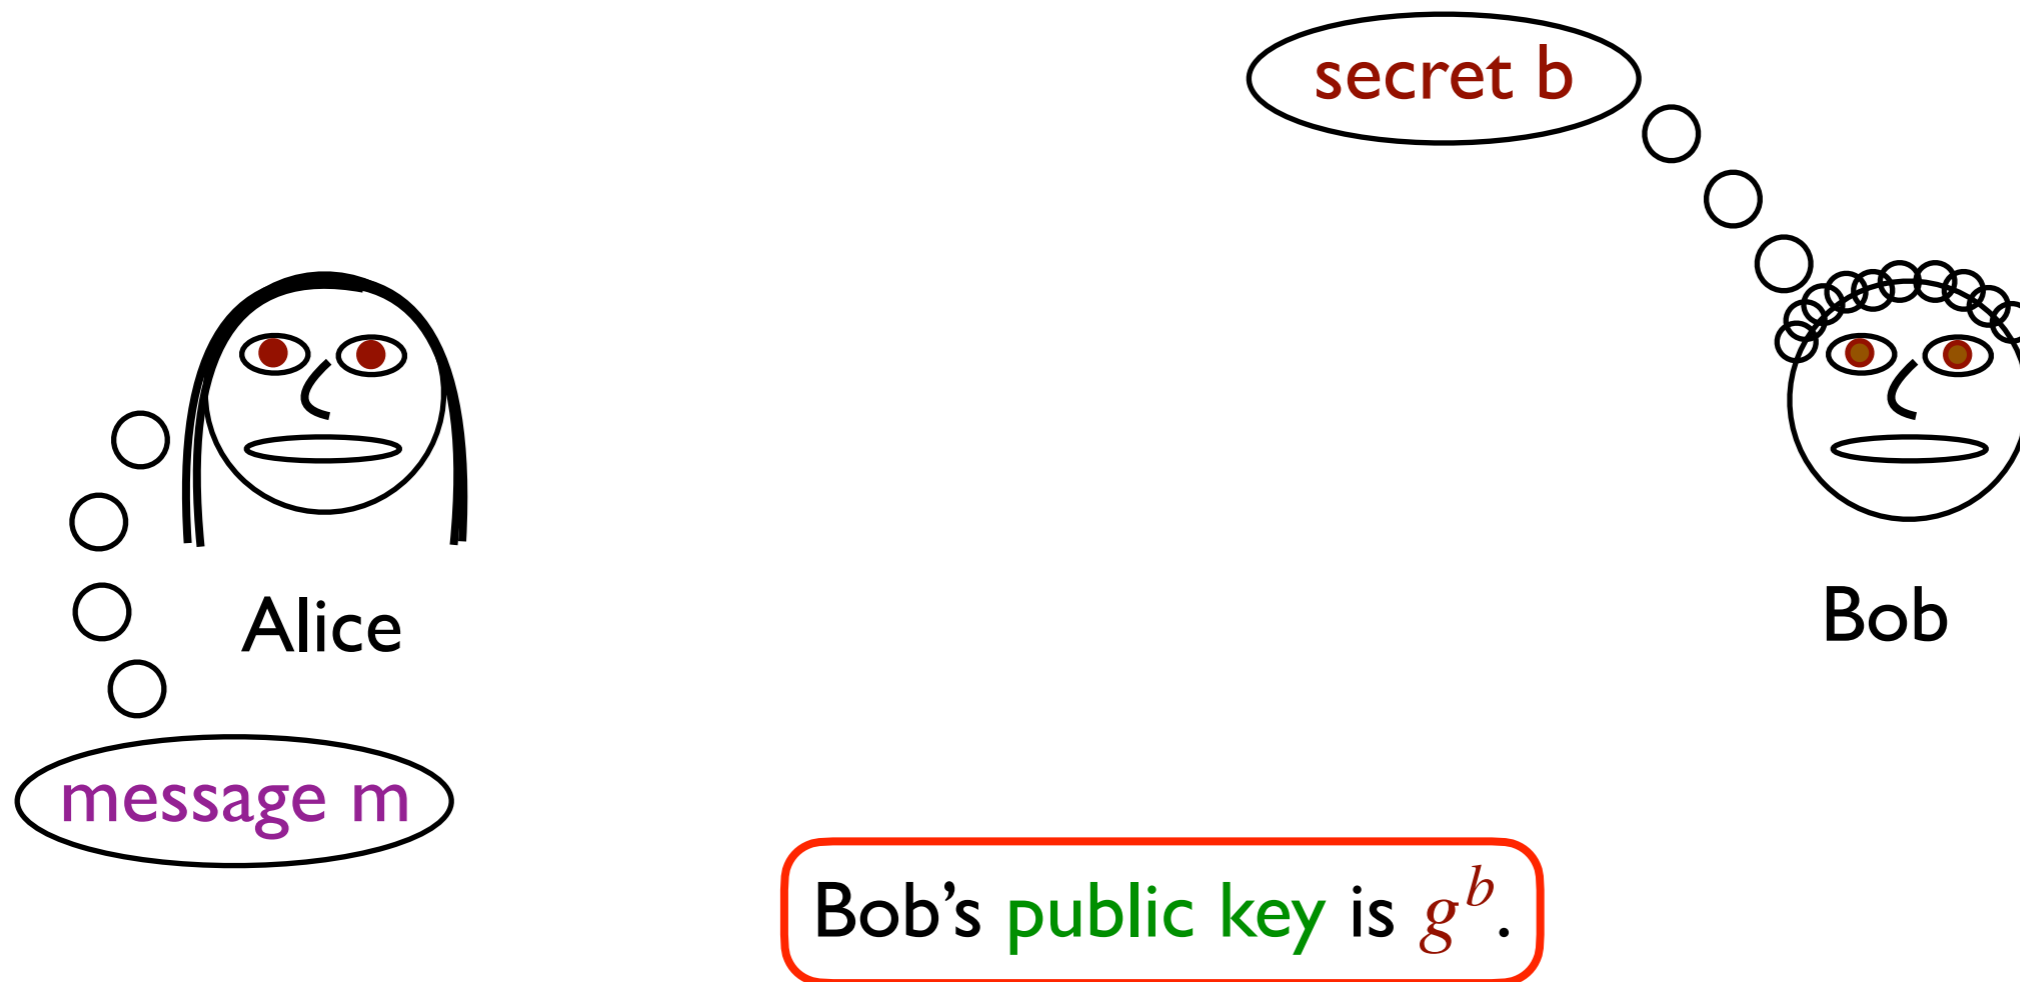
- In class
- Open book (including textbook), no electronic devices
- Will cover material through Diffie-Hellman and key exchange, but not public key encryption.

There will be a poll on Piazza out tonight on which topics to review on Tuesday, Oct. 17. Vote by Sunday night.

The solutions to PS #5 are available on ELMS, although the problem set is not fully graded. I have also put a copy of last year's midterm exam on ELMS if you want to practice. There is also a list of selected practice problems from the textbook available on the course web page.

El Gamal Encryption


When Alice sends g^a , she can also send an encrypted message m , for instance with ciphertext $c = m \cdot g^{ab}$.



Advantage over Diffie-Hellman: **non-interactive**

El Gamal Encryption

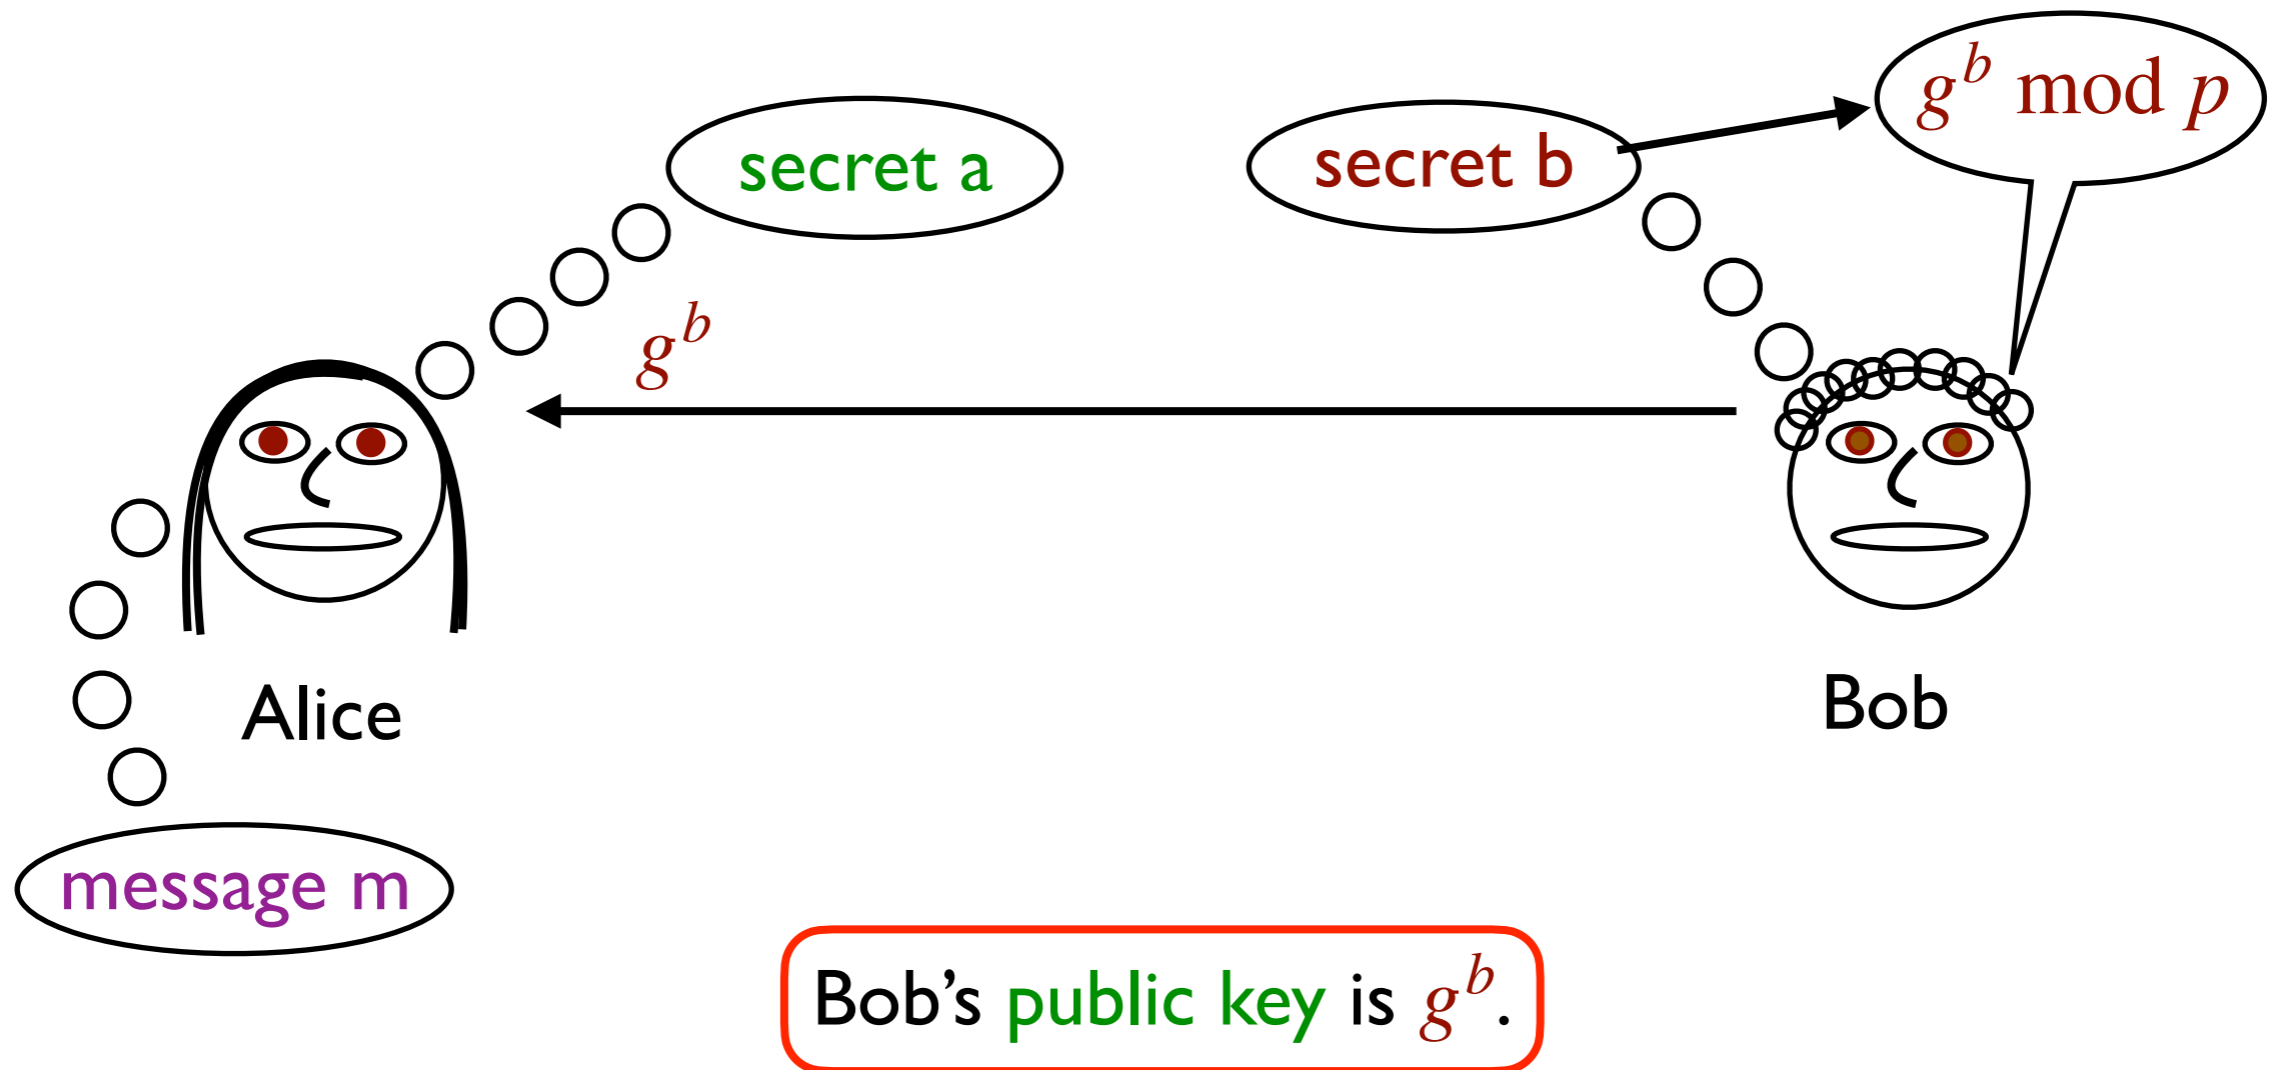
When Alice sends g^a , she can also send an encrypted message m , for instance with ciphertext $c = m \cdot g^{ab}$.



Advantage over Diffie-Hellman: **non-interactive**

El Gamal Encryption


When Alice sends g^a , she can also send an encrypted message m , for instance with ciphertext $c = m \cdot g^{ab}$.



Advantage over Diffie-Hellman: **non-interactive**

El Gamal Encryption


When Alice sends g^a , she can also send an encrypted message m , for instance with ciphertext $c = m \cdot g^{ab}$.



Advantage over Diffie-Hellman: **non-interactive**

El Gamal Encryption

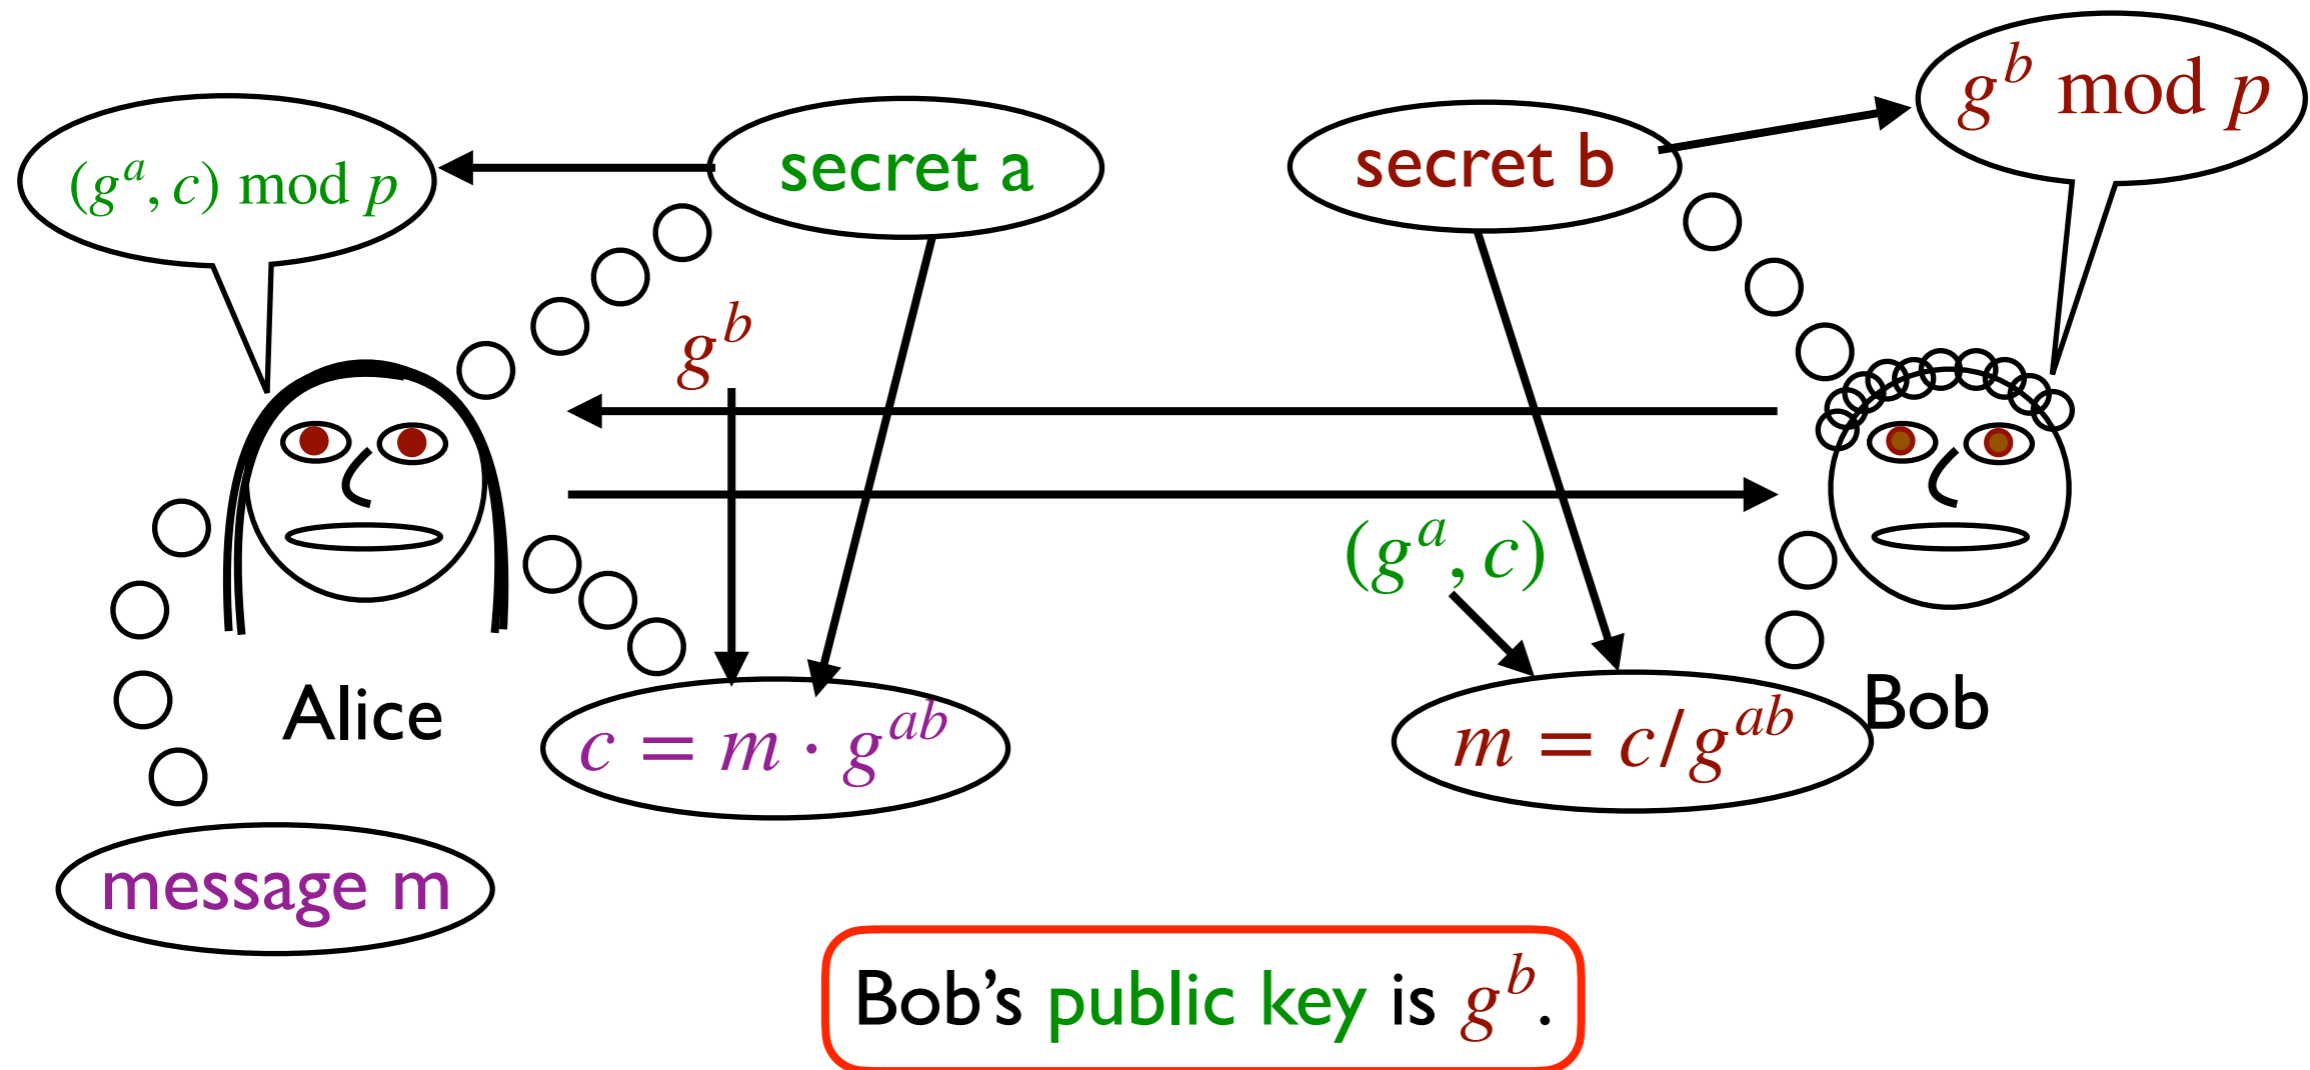
When Alice sends g^a , she can also send an encrypted message m , for instance with ciphertext $c = m \cdot g^{ab}$.



Advantage over Diffie-Hellman: **non-interactive**


El Gamal Encryption

When Alice sends g^a , she can also send an encrypted message m , for instance with ciphertext $c = m \cdot g^{ab}$.




Advantage over Diffie-Hellman: **non-interactive**


Public Key Encryption



Alice




Bob



Eve


Public-key encryption is an **asymmetric** protocol.

Public Key Encryption




Public-key encryption is an **asymmetric** protocol.

Public Key Encryption




Public-key encryption is an **asymmetric** protocol.

Public Key Encryption




Public-key encryption is an **asymmetric** protocol.

Public Key Encryption




Public-key encryption is an **asymmetric** protocol.

Public Key Encryption



Public-key encryption is an **asymmetric** protocol.

Public Key Encryption



Public-key encryption is an **asymmetric** protocol.

Definition of Public Key Encryption

Definition: A **public-key encryption protocol** is a set of three probabilistic polynomial-time algorithms (**Gen**, **Enc**, **Dec**).

Gen is the **key generation algorithm**. It takes as input **s**, the **security parameter**, and outputs a **public key, private key pair** $(e, d) \in \{0,1\}^* \times \{0,1\}^*$.

Enc is the **encryption algorithm**. It takes as input **e** and a **plaintext or message** $m \in \{0,1\}^*$ and outputs a **ciphertext** $c \in \{0,1\}^*$.

Dec is the **decryption algorithm**. It takes as input **d** and **c** and outputs some $m' \in \{0,1\}^*$.

The encryption protocol is **correct** if

$$Dec(d, Enc(e, m)) = m$$

Note: **Gen** here is much more complex than for private-key encryption and doesn't just generate random bit strings.

El Gamal as a Public Key Scheme

Gen: The group G and base g have been pre-determined. **Gen** chooses a random $b \in \{0, \dots, \text{ord}(g) - 1\}$, which becomes the private key. I.e., $d = b$. The public key is $e = g^b$.

Enc: Given message m and public key e . Choose random $a \in \{0, \dots, \text{ord}(g) - 1\}$. The ciphertext is $c = (g^a, m \cdot e^a)$.


Dec: Given $c = (A, x)$ and d . The decrypted message is $m' = x/A^d$.

G and g can instead be generated using **Gen** and made part of the public key.

(All calculations here are done within the group G .)

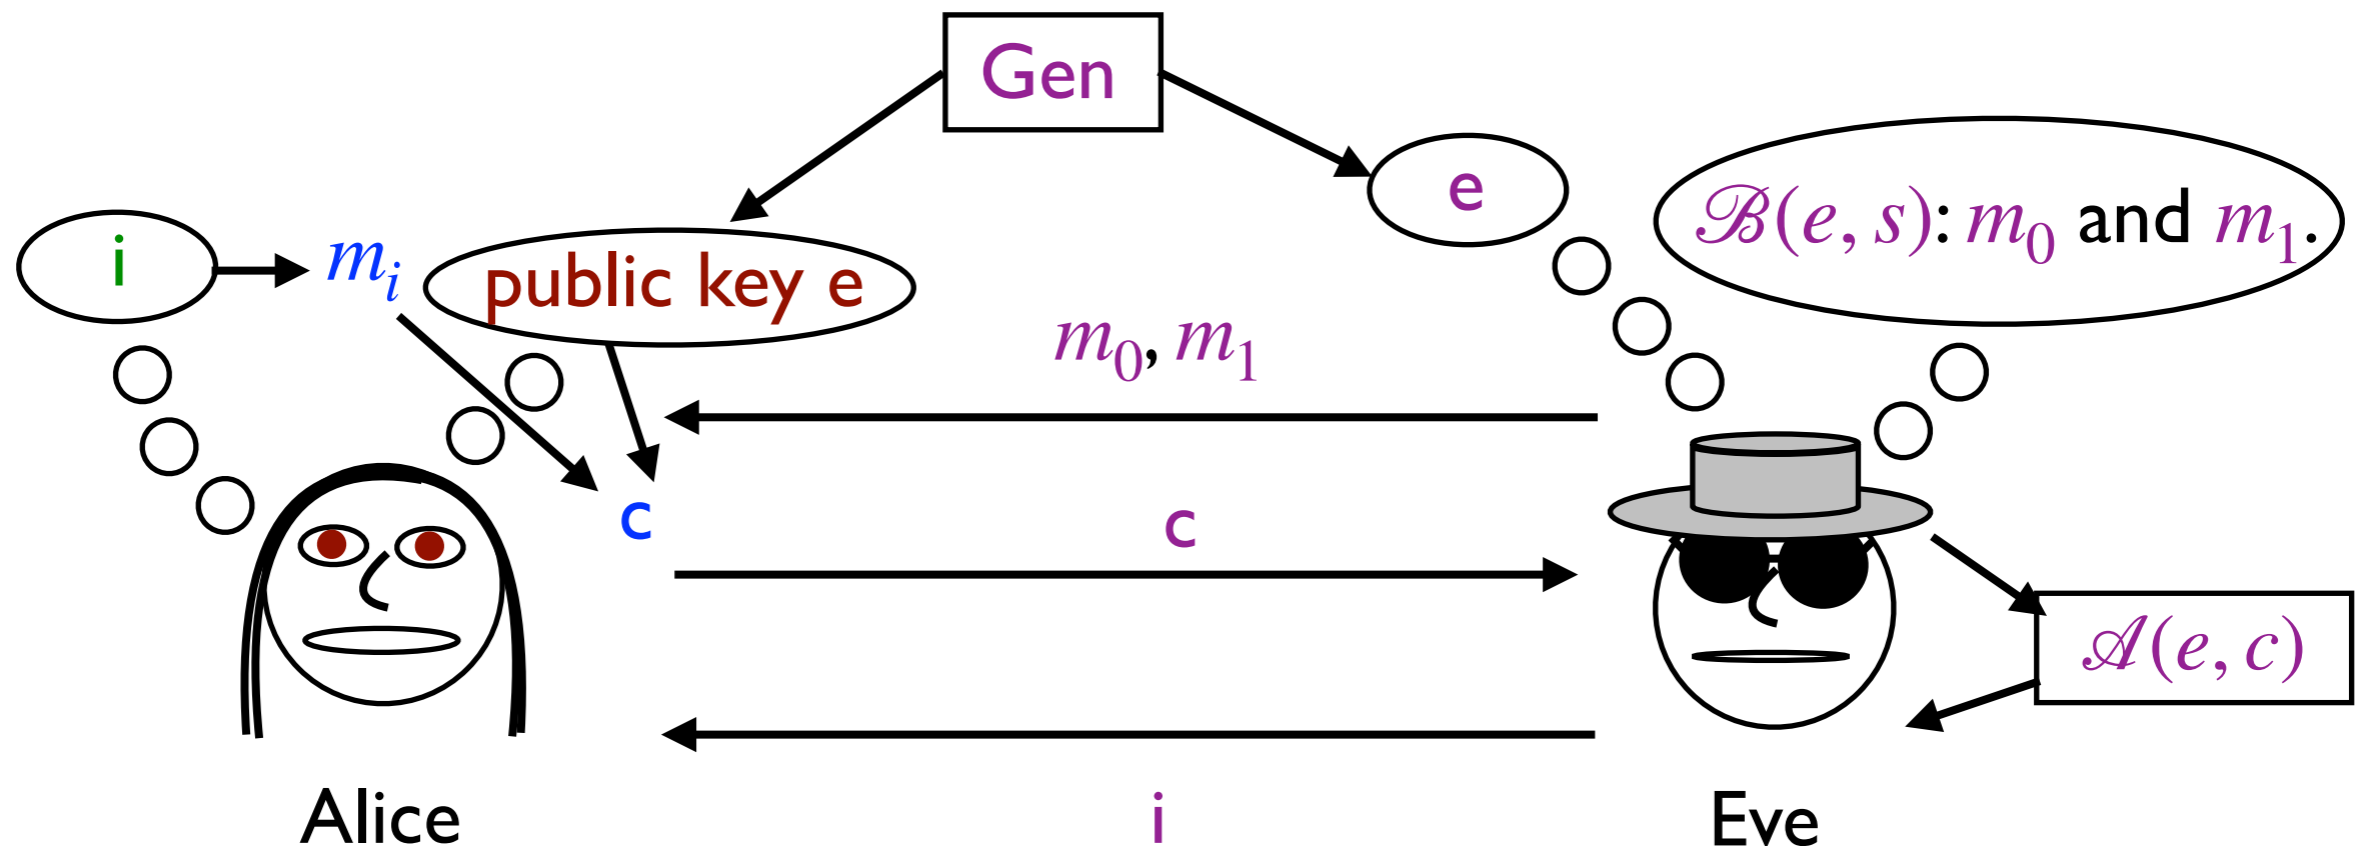
Security of Public Key Encryption

The definition of EAV security for public key encryption is very similar to EAV security for private key encryption.



Security of Public Key Encryption

The definition of EAV security for public key encryption is very similar to EAV security for private key encryption.



The main difference is that now Eve is **given the public key e** . She can use e to choose the messages she wants to use as challenges as well as in her attack to decipher the message.

Public Key EAV Security

Definition: A public-key encryption protocol $(\text{Gen}, \text{Enc}, \text{Dec})$ with security parameter s has **indistinguishable encryptions in the presence of an eavesdropper (is EAV-secure)** if, for any pair of messages m_0 and m_1 chosen by the adversary (using efficient algorithm $\mathcal{B}(e, s)$) and for any efficient attack $\mathcal{A}(e, c)$,

$$| \Pr_{(e,d)}(\mathcal{A}(e, \text{Enc}(e, m_0)) = 1) - \Pr_{(e,d)}(\mathcal{A}(e, \text{Enc}(e, m_1)) = 1) | \leq \epsilon(s)$$

for negligible $\epsilon(s)$ and probability taken over valid public key, private key pairs (e,d) and randomness of Enc , \mathcal{A} , and \mathcal{B} .

Public Key EAV Security

Definition: A public-key encryption protocol $(\text{Gen}, \text{Enc}, \text{Dec})$ with security parameter s has **indistinguishable encryptions in the presence of an eavesdropper (is EAV-secure)** if, for any pair of messages m_0 and m_1 chosen by the adversary (using efficient algorithm $\mathcal{B}(e, s)$) and for any efficient attack $\mathcal{A}(e, c)$,

$$| \Pr_{(e,d)}(\mathcal{A}(e, \text{Enc}(e, m_0)) = 1) - \Pr_{(e,d)}(\mathcal{A}(e, \text{Enc}(e, m_1)) = 1) | \leq \epsilon(s)$$

for negligible $\epsilon(s)$ and probability taken over valid public key, private key pairs (e,d) and randomness of Enc , \mathcal{A} , and \mathcal{B} .

Note that d and Dec are not needed in this definition. But they are still important because the protocol cannot be **correct** if they are not chosen right.

CPA Security with Public Keys

Suggestions: Does anyone have an idea how to modify this definition to get CPA security?

CPA Security with Public Keys

Suggestions: Does anyone have an idea how to modify this definition to get CPA security?

How about we don't change anything?

CPA Security with Public Keys

Suggestions: Does anyone have an idea how to modify this definition to get CPA security?

How about we don't change anything?

For private key cryptography, to get CPA security, we modified the definition of EAV security by giving Eve access to an oracle that performed $\text{Enc}(k,x)$. This enables Eve to create plaintext - ciphertext pairs as she wishes.

CPA Security with Public Keys

Suggestions: Does anyone have an idea how to modify this definition to get CPA security?

How about we don't change anything?

For private key cryptography, to get CPA security, we modified the definition of EAV security by giving Eve access to an oracle that performed $\text{Enc}(k,x)$. This enables Eve to create plaintext - ciphertext pairs as she wishes.

But Eve has the public key. She doesn't *need* an oracle — she can already create plaintext - ciphertext pairs using the public key.

CPA Security with Public Keys

Suggestions: Does anyone have an idea how to modify this definition to get CPA security?

How about we don't change anything?

For private key cryptography, to get CPA security, we modified the definition of EAV security by giving Eve access to an oracle that performed $\text{Enc}(k,x)$. This enables Eve to create plaintext - ciphertext pairs as she wishes.

But Eve has the public key. She doesn't *need* an oracle — she can already create plaintext - ciphertext pairs using the public key.

For public key encryption, **EAV security is the same as CPA security!**

Key Encapsulation

Both of the techniques I mentioned so far for encryption (El Gamal and using the key directly in the pseudo one-time pad) suffer from being very slow and inefficient.


But what if we use the key generated by Diffie-Hellman for a symmetric cryptosystem? Then we only have to use Diffie-Hellman for the initial set-up and the rest of the message can be sent with the more efficient symmetric encryption protocol.

We want to keep the non-interactive part of El Gamal, but we don't need (or want) to encrypt a message with it. We do, however, need to convert the key from a group element to a bit string using a key derivation function.


We will get a procedure to create a shared random encrypted key using a public key. This is known as a **key encapsulation mechanism (KEM)**.

KEM with Diffie-Hellman/El Gamal

Fixed as part of the protocol: p, g, H . Bob's public key is g^b .



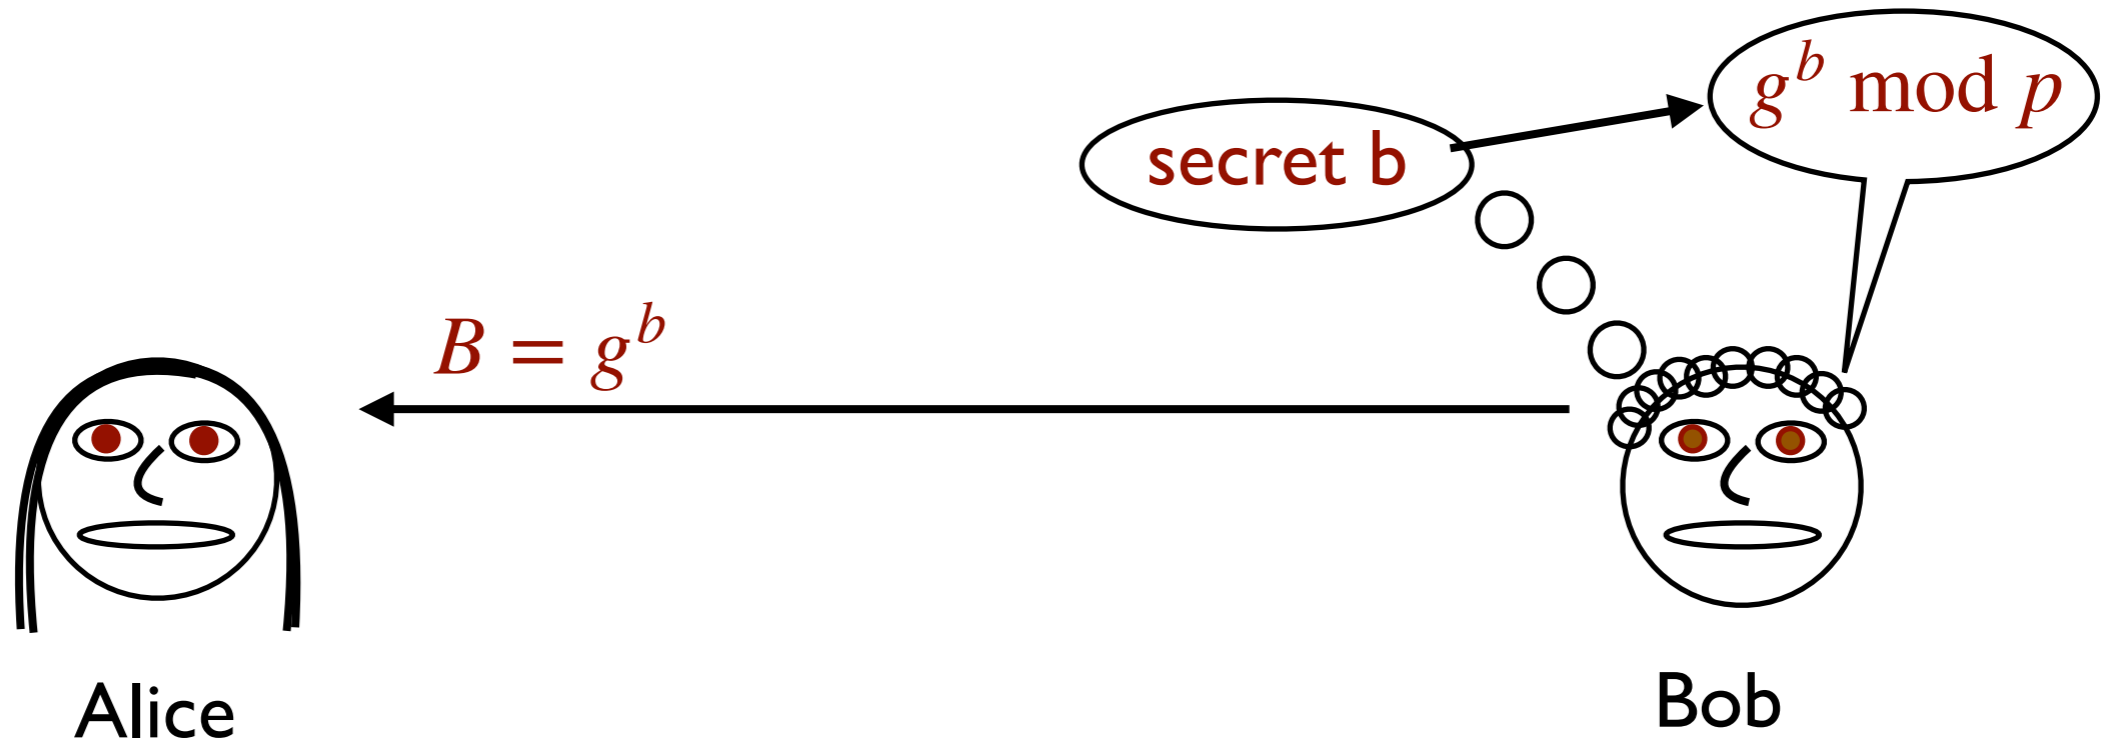
Alice



Bob

KEM with Diffie-Hellman/El Gamal

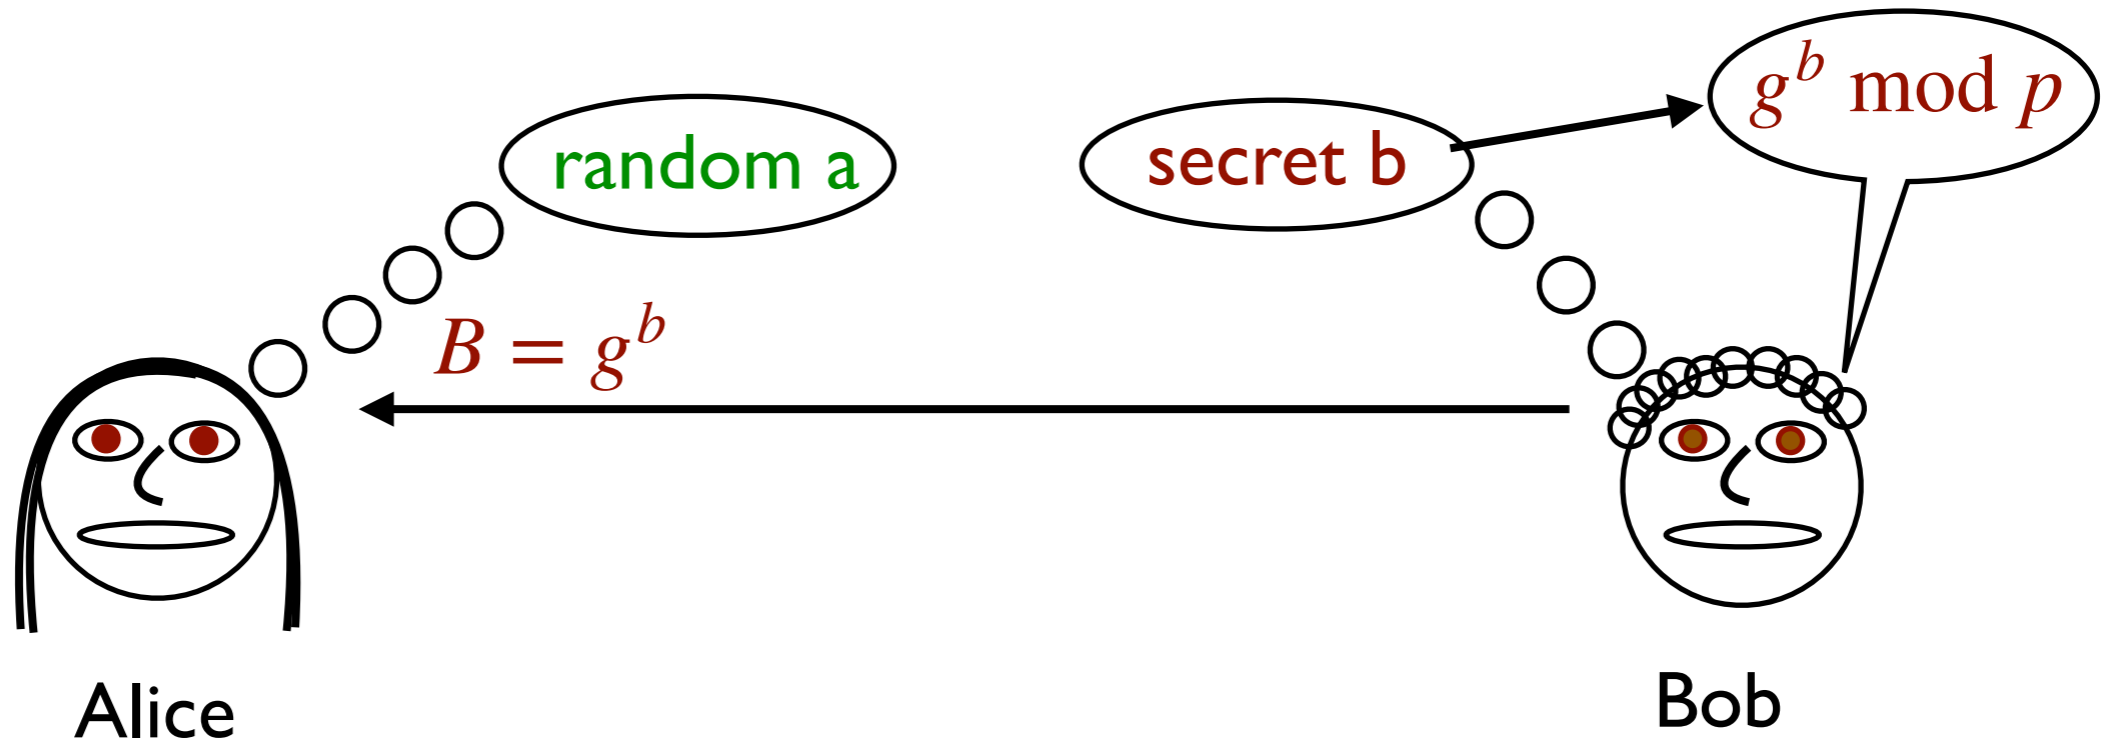
Fixed as part of the protocol: p, g, H . Bob's public key is g^b .



I. First, Bob distributes his public key B .

KEM with Diffie-Hellman/El Gamal

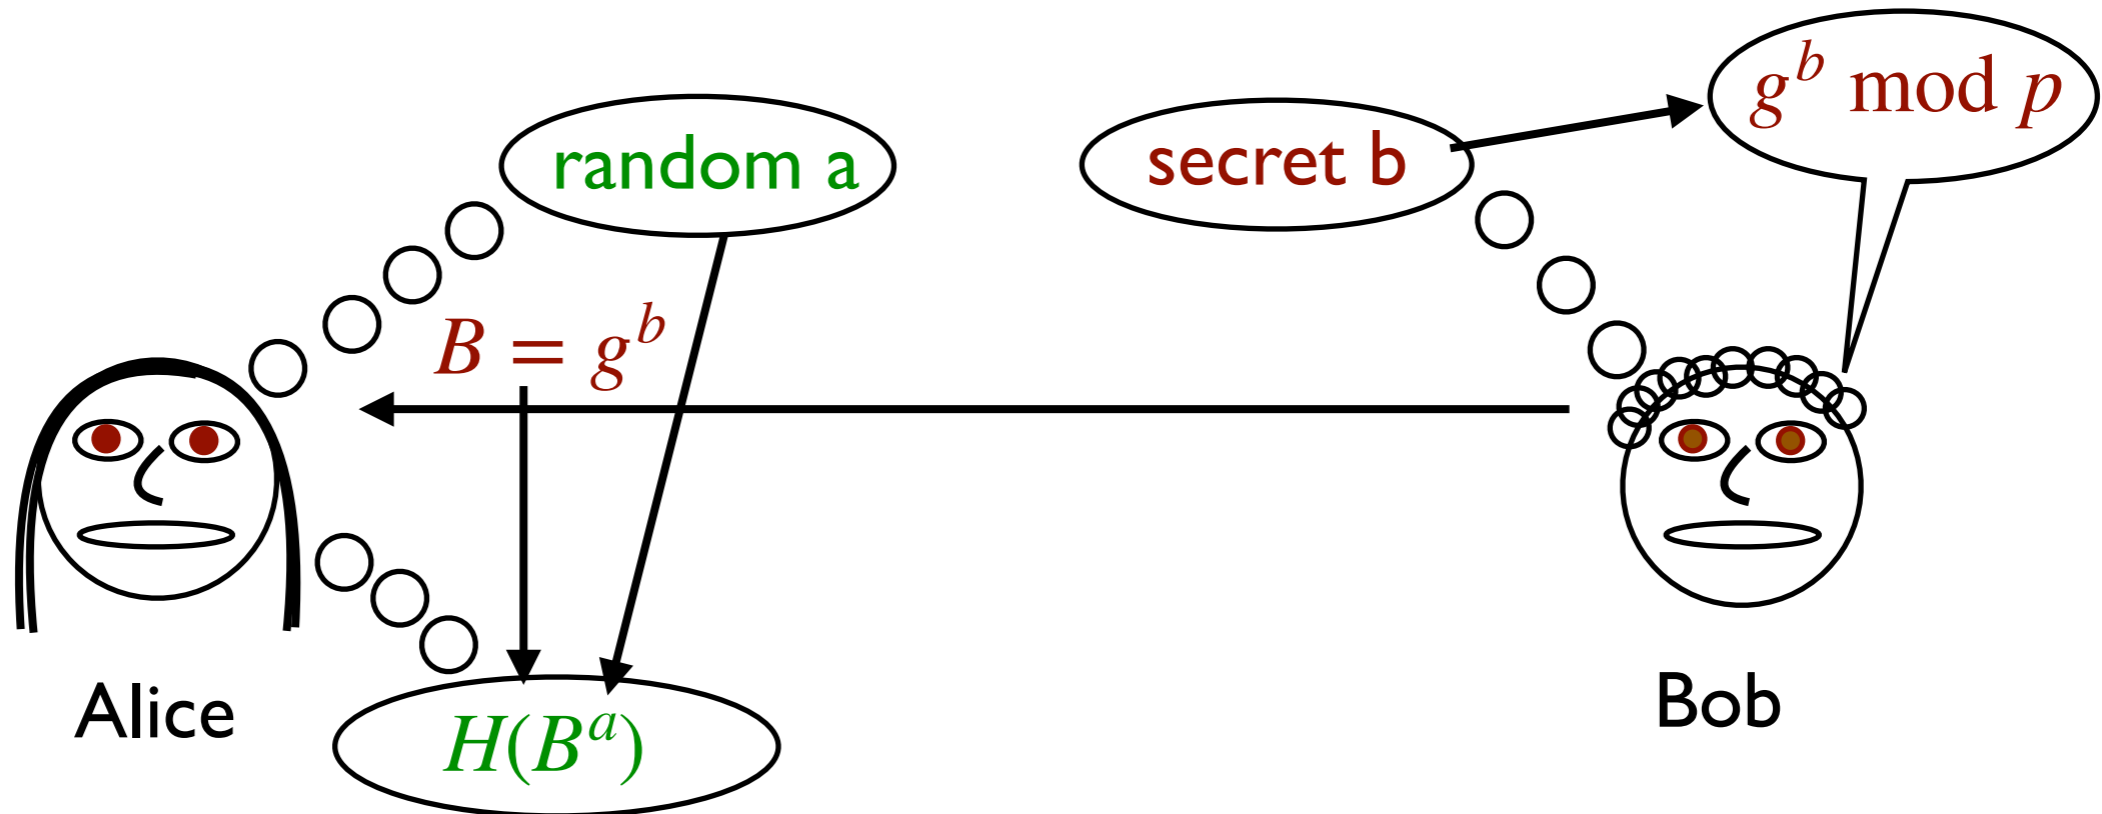
Fixed as part of the protocol: p, g, H . Bob's public key is g^b .



I. First, Bob distributes his public key B .

KEM with Diffie-Hellman/El Gamal


Fixed as part of the protocol: p, g, H . Bob's public key is g^b .



1. First, Bob distributes his public key B .
2. Alice generates a random key $H(g^{ab})$

KEM with Diffie-Hellman/El Gamal

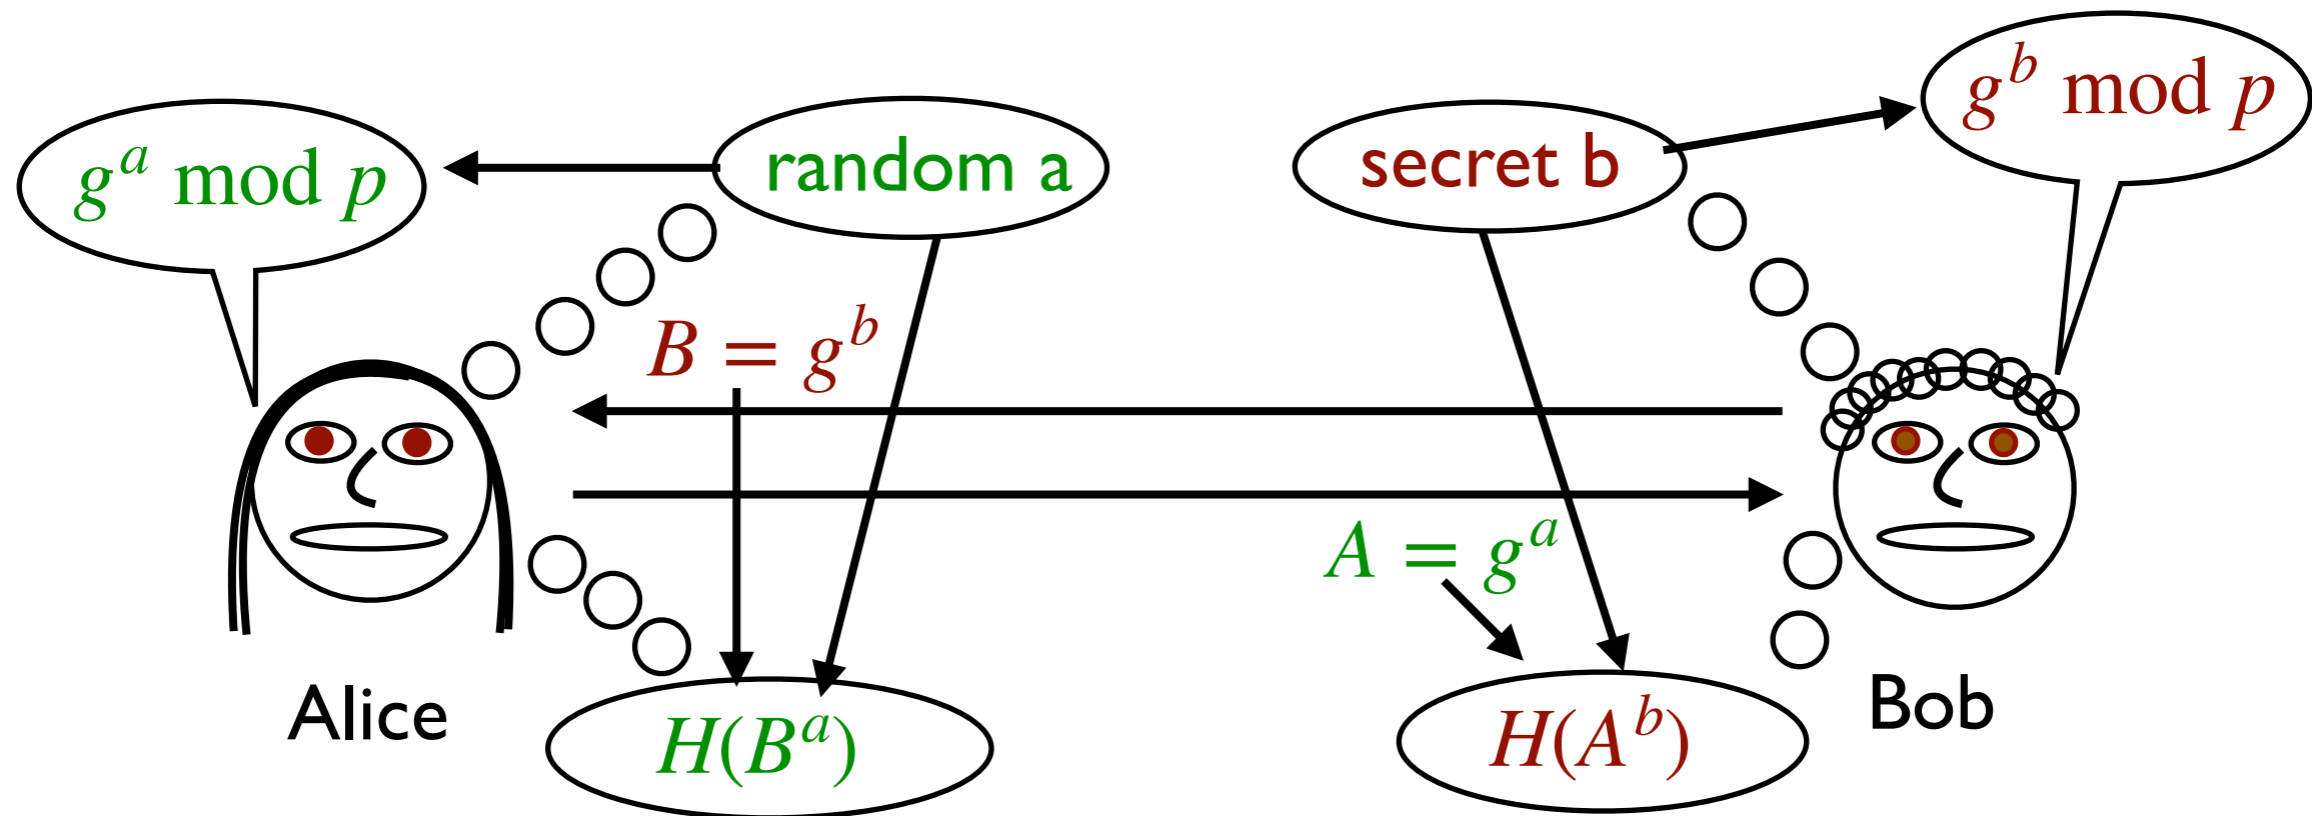
Fixed as part of the protocol: p, g, H . Bob's public key is g^b .



1. First, Bob distributes his public key B .
2. Alice generates a random key $H(g^{ab})$
3. Alice sends A to Bob as an encrypted way of transmitting the key.

KEM with Diffie-Hellman/El Gamal

Fixed as part of the protocol: p, g, H . Bob's public key is g^b .



1. First, Bob distributes his public key B .
2. Alice generates a random key $H(g^{ab})$
3. Alice sends A to Bob as an encrypted way of transmitting the key.
4. Bob decrypts the key using his private key to get $H(g^{ab})$.

Definition of Key Encapsulation

Definition: A **key encapsulation mechanism** is a set of three probabilistic polynomial-time algorithms (**Gen**, **Encaps**, **Decaps**).

Gen is the **key generation algorithm**. It takes as input **s**, the **security parameter**, and outputs a **public key, private key pair** $(e, d) \in \{0,1\}^* \times \{0,1\}^*$.

Encaps is the **encapsulation algorithm**. It takes as input **e** (only) and outputs a **ciphertext** $c \in \{0,1\}^*$ and a key $k \in \{0,1\}^{\ell(s)}$, for some function $\ell(s)$ (the **key length**).

Decaps is the **decapsulation algorithm**. It takes as input **d** and **c** and outputs some $k' \in \{0,1\}^{\ell(s)}$.

If **Encaps(e)** outputs ciphertext **c** and key **k**, then we require that **Decaps(d,c)** outputs $k' = k$.

Diffie-Hellman/El Gamal KEM

Gen: The group G , base g , and key derivation function $H(x)$ have been pre-determined. **Gen** chooses a random $b \in \{0, \dots, \text{ord}(g) - 1\}$, which becomes the private key. I.e., $d = b$. The public key is $e = g^b$.

Encaps: Choose random $a \in \{0, \dots, \text{ord}(g) - 1\}$. The ciphertext is $c = g^a$ and the key is $k = H(e^a)$.


Decaps: Given c and d . The key is $k' = H(c^d)$.

(Again, all calculations are done within the group G .)


KEM/DEM

A KEM creates and sends a random key string to someone whose public key you have.

That key then becomes the key for a private-key encryption system, which is called a **data-encapsulation mechanism (DEM)** in this context.



Alice




Bob

KEM/DEM

A KEM creates and sends a random key string to someone whose public key you have.


That key then becomes the key for a private-key encryption system, which is called a **data-encapsulation mechanism (DEM)** in this context.



KEM/DEM

A KEM creates and sends a random key string to someone whose public key you have.


That key then becomes the key for a private-key encryption system, which is called a **data-encapsulation mechanism (DEM)** in this context.



KEM/DEM

A KEM creates and sends a random key string to someone whose public key you have.

That key then becomes the key for a private-key encryption system, which is called a **data-encapsulation mechanism (DEM)** in this context.

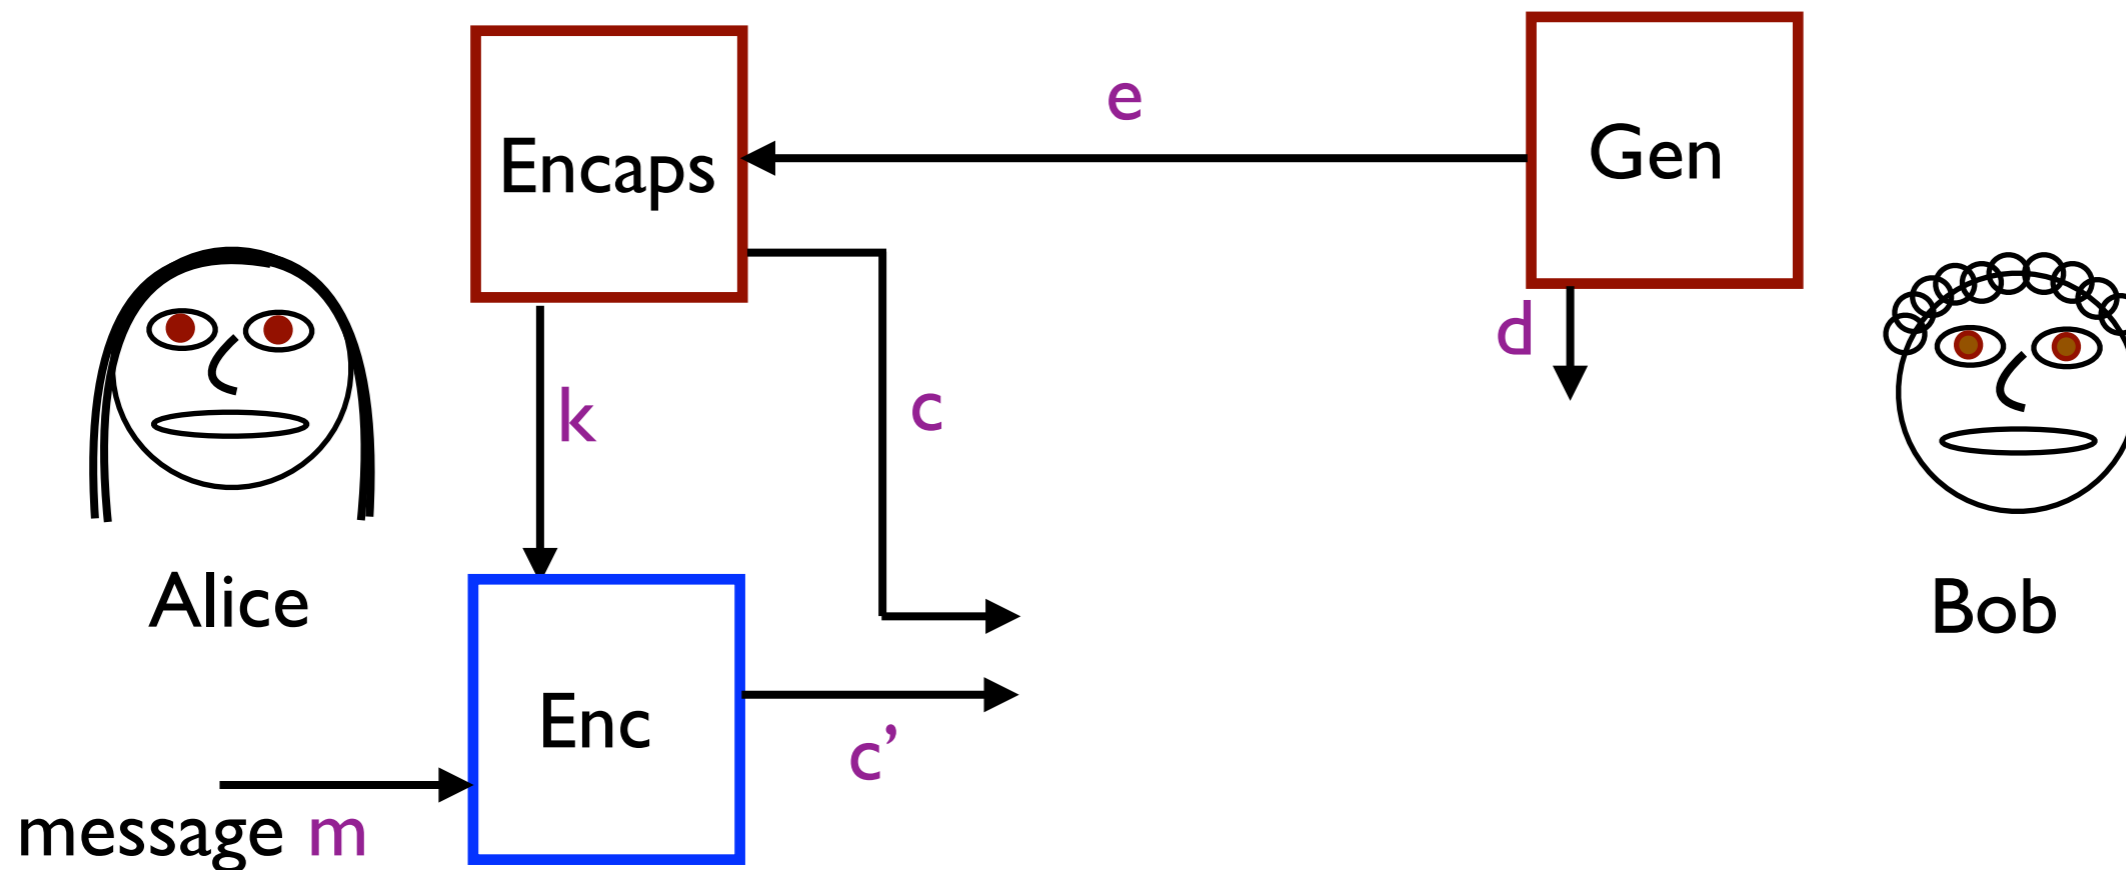


message m

KEM/DEM

A KEM creates and sends a random key string to someone whose public key you have.

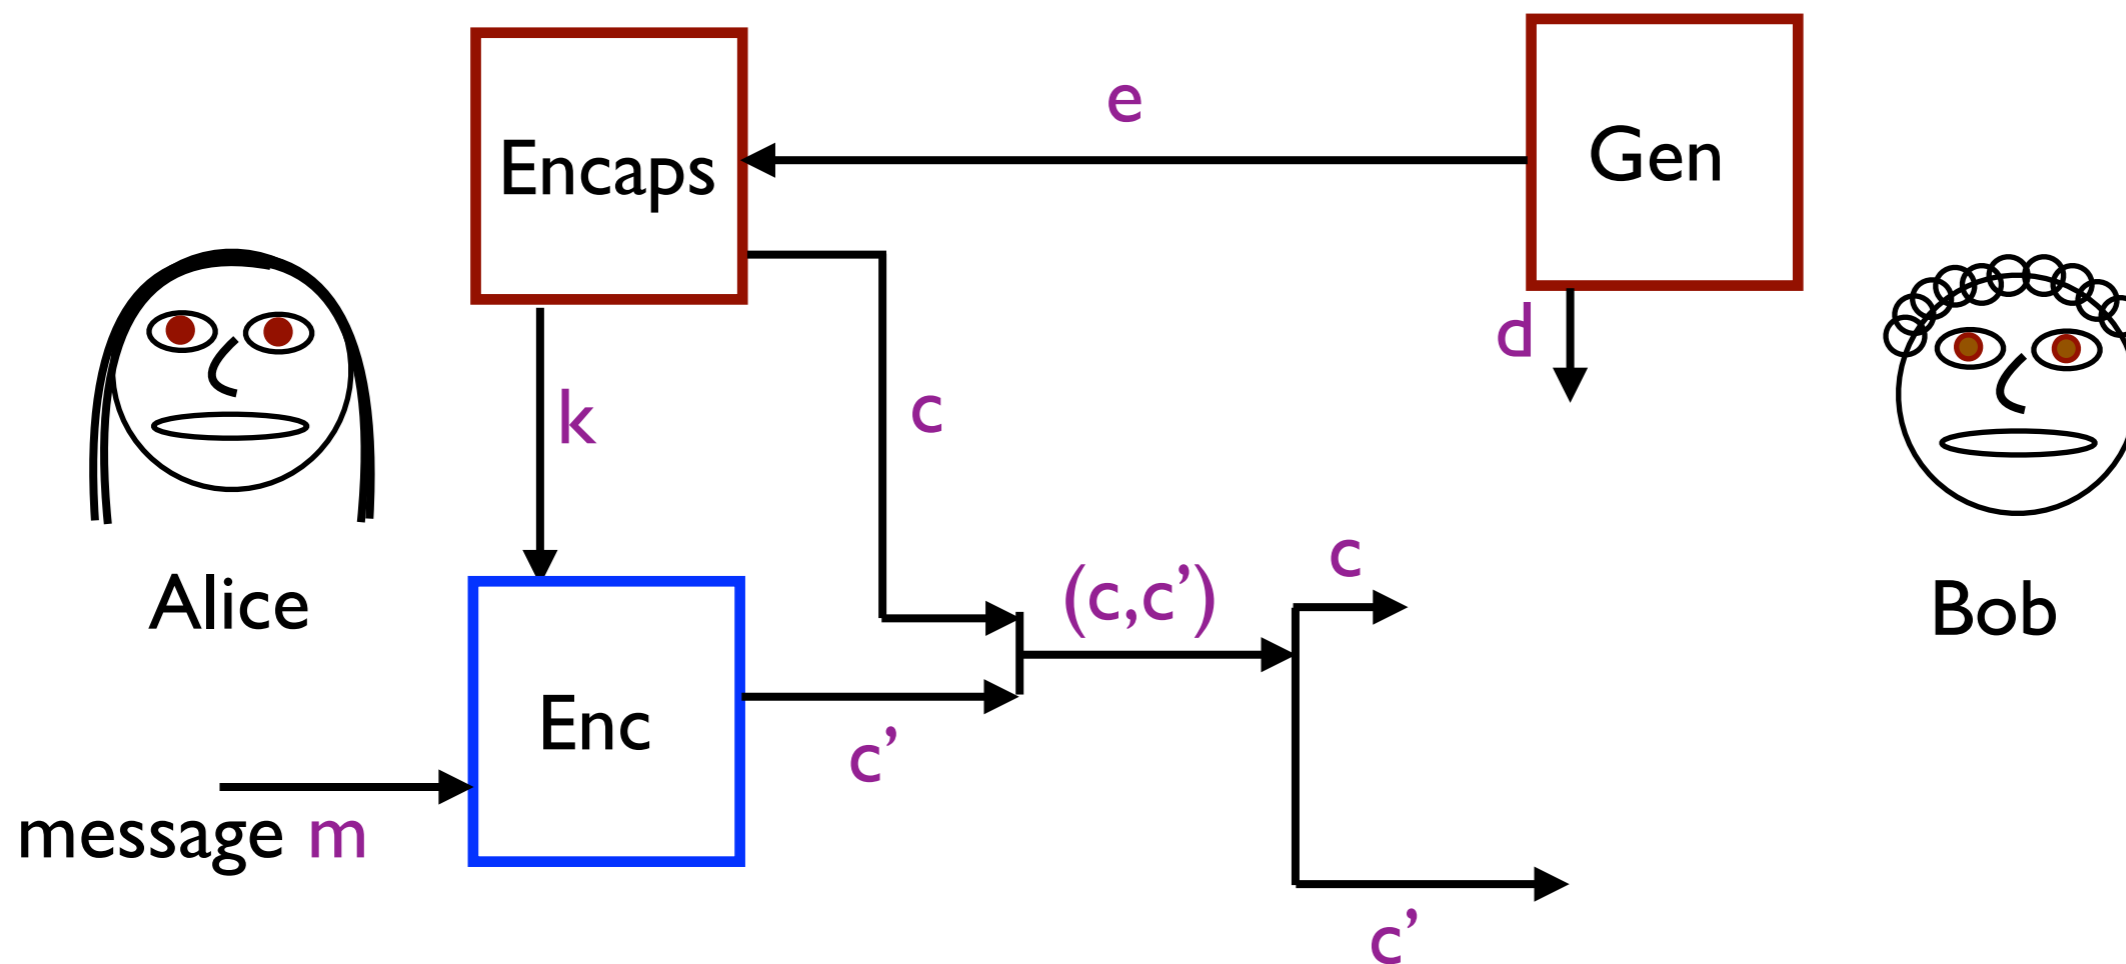
That key then becomes the key for a private-key encryption system, which is called a **data-encapsulation mechanism (DEM)** in this context.



KEM/DEM

A KEM creates and sends a random key string to someone whose public key you have.

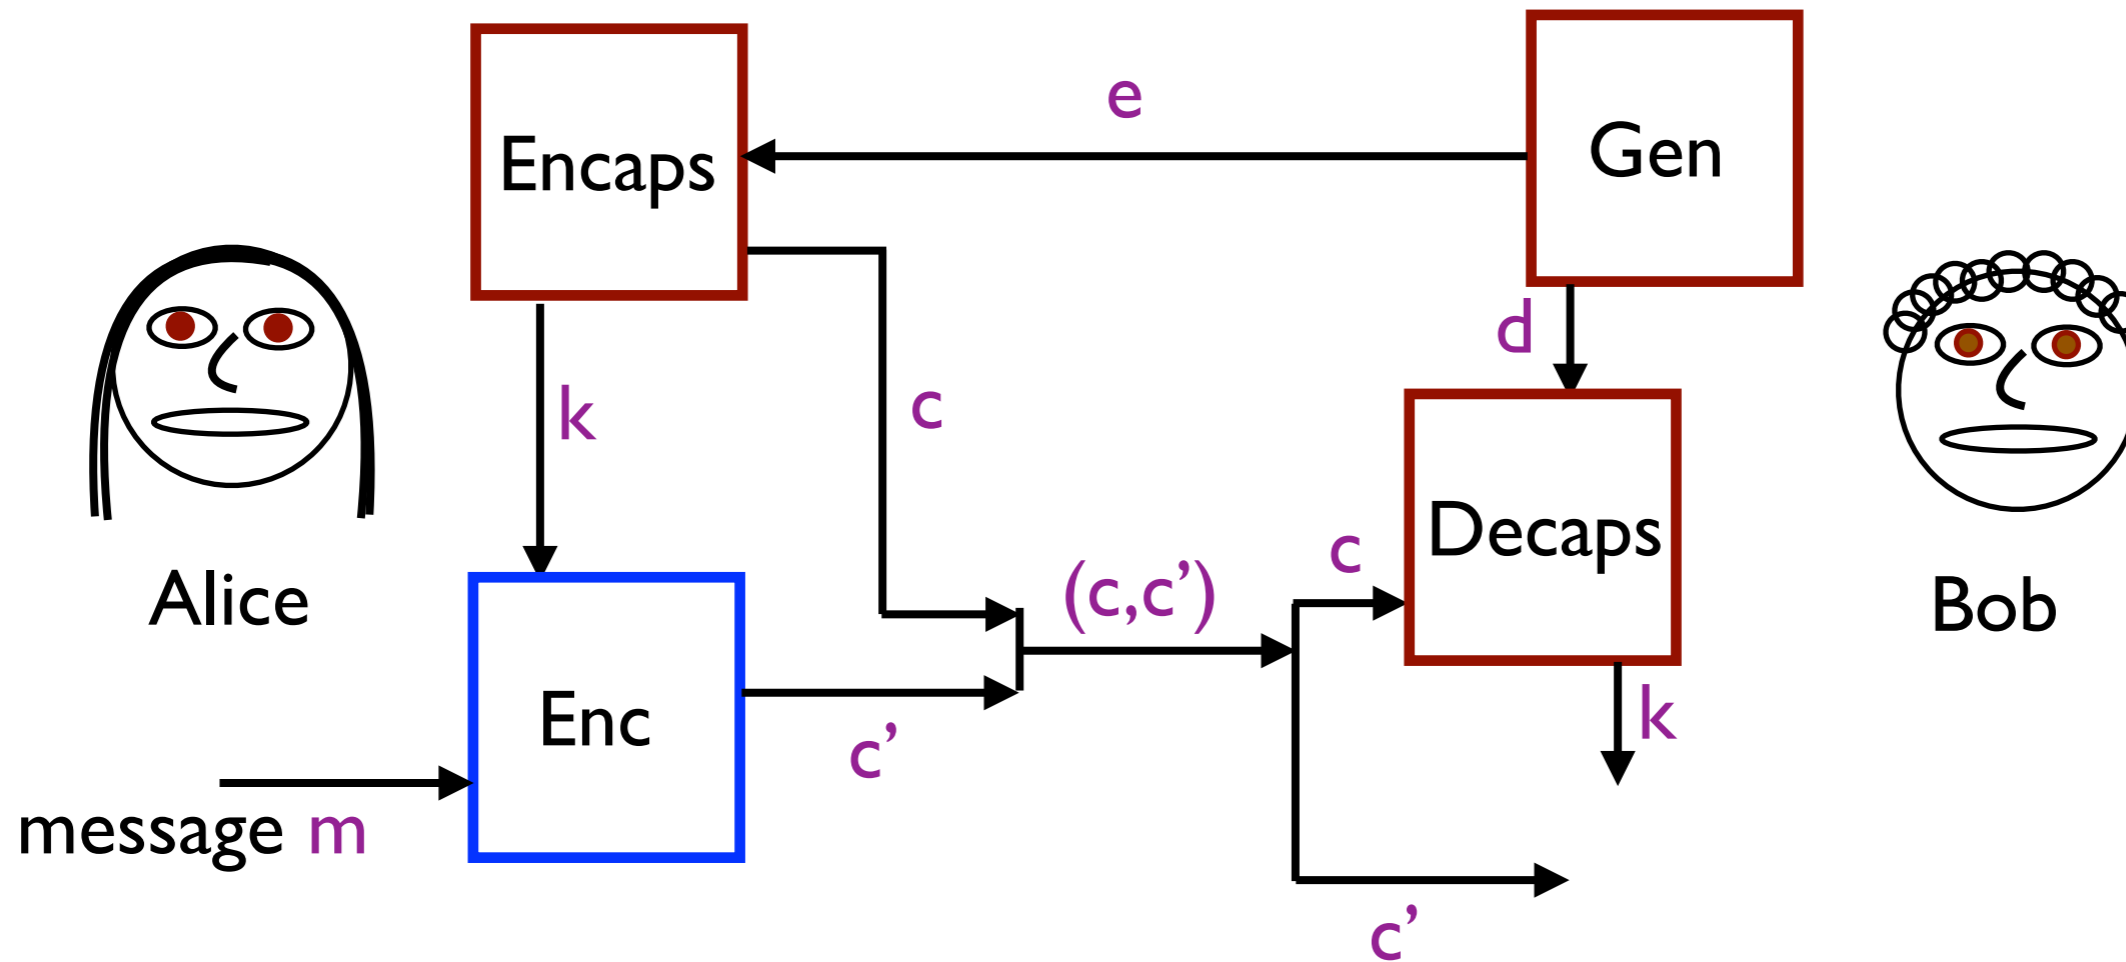
That key then becomes the key for a private-key encryption system, which is called a **data-encapsulation mechanism (DEM)** in this context.



KEM/DEM

A KEM creates and sends a random key string to someone whose public key you have.

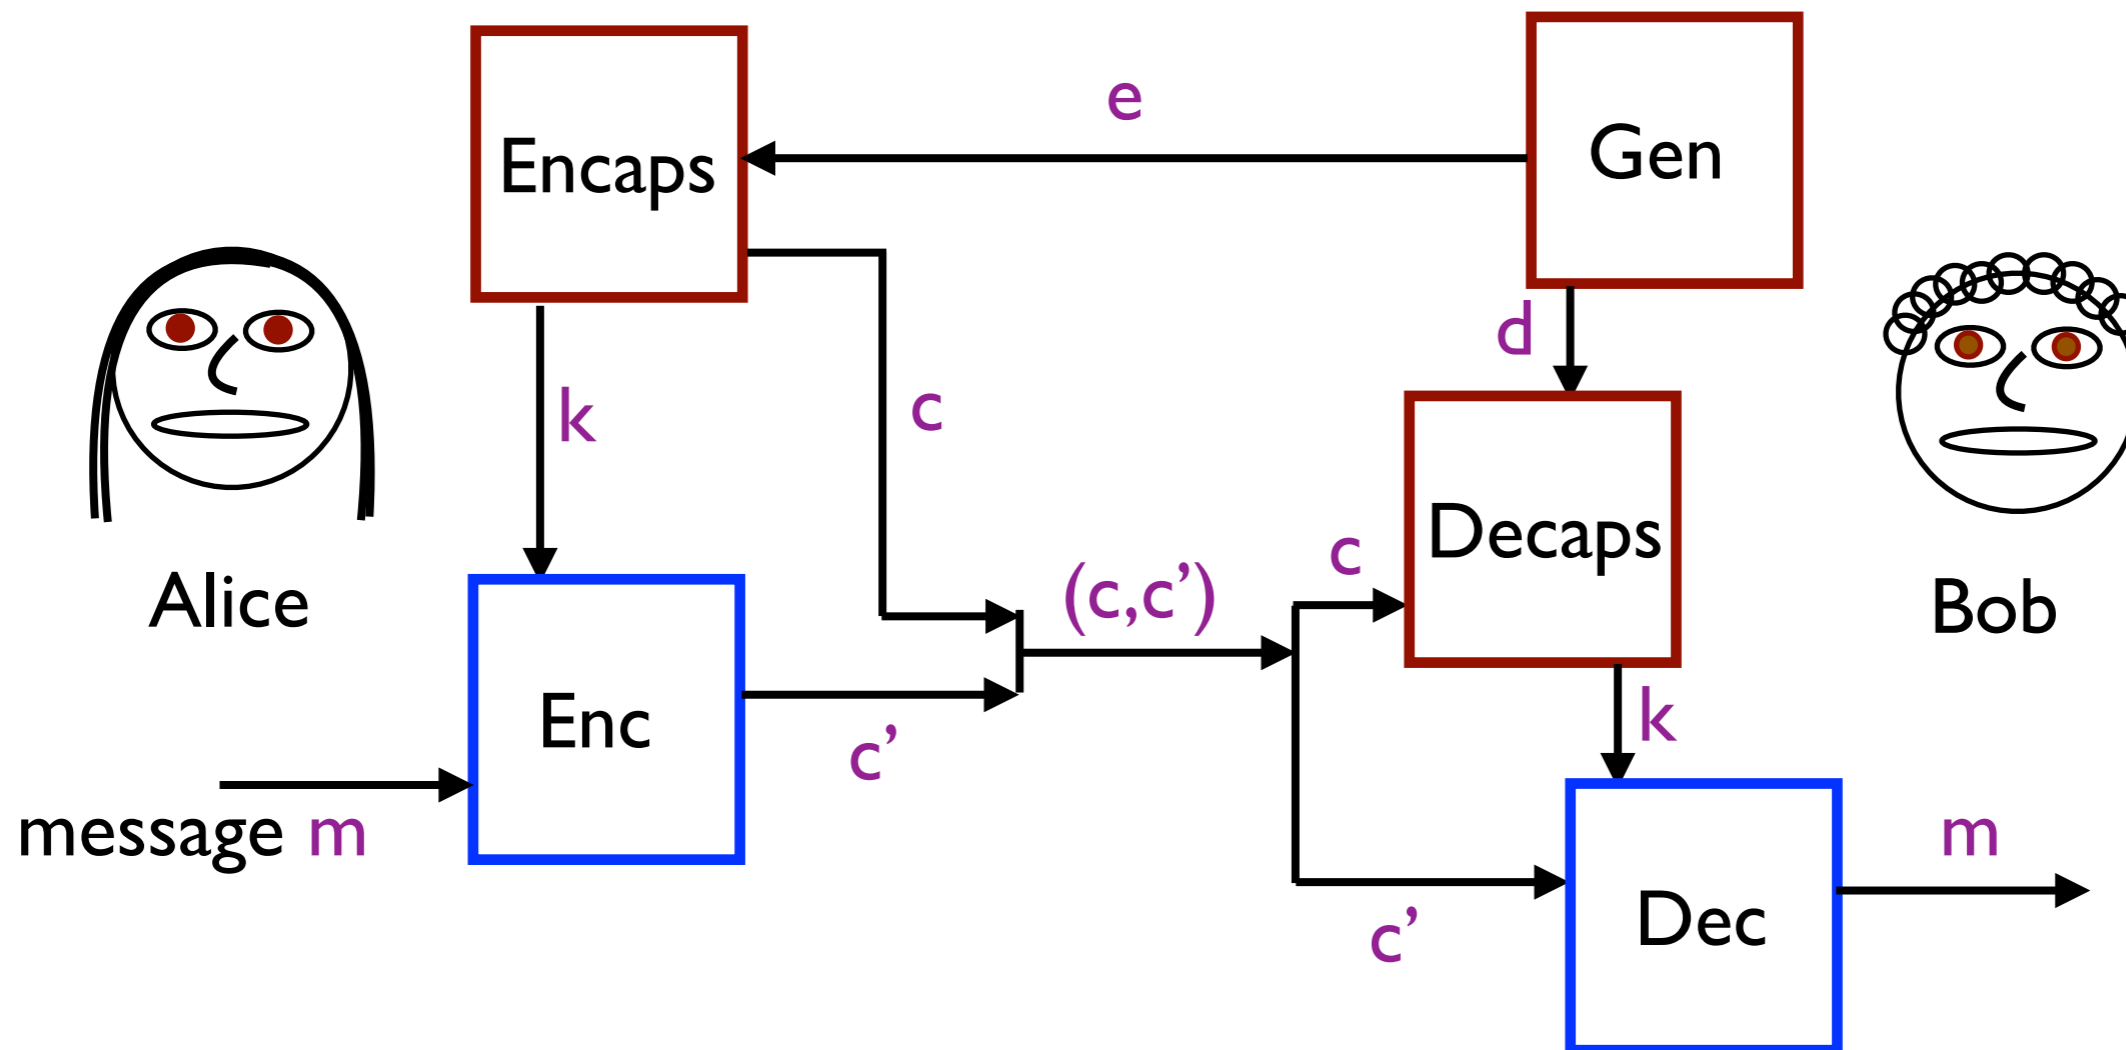
That key then becomes the key for a private-key encryption system, which is called a **data-encapsulation mechanism (DEM)** in this context.



KEM/DEM

A KEM creates and sends a random key string to someone whose public key you have.

That key then becomes the key for a private-key encryption system, which is called a **data-encapsulation mechanism (DEM)** in this context.



Why KEM/DEM?

Public key cryptosystems are much slower, more inefficient in terms of communication, and less secure than private key cryptosystems.

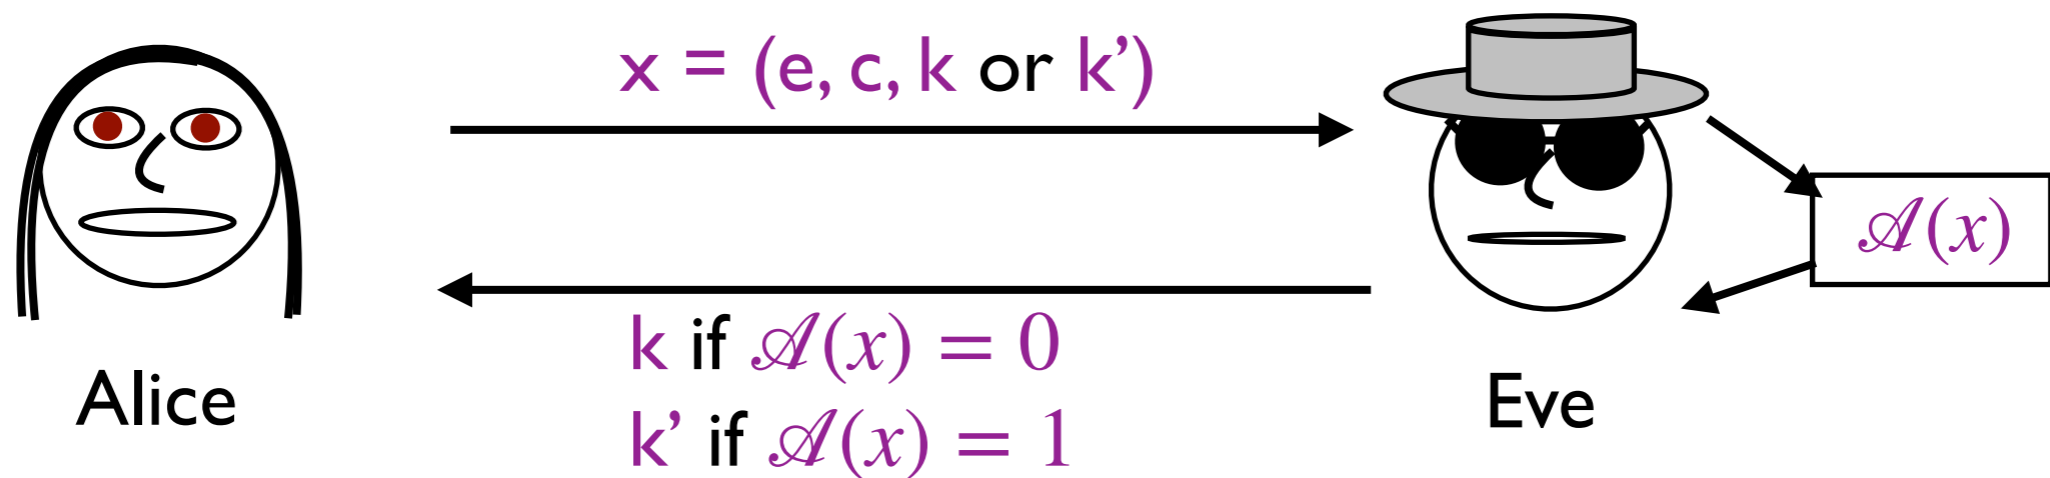
Private key cryptosystems require a previously shared key, making it hard to communicate with someone new and making key management a headache with many people.

KEM/DEM combines the efficiency of private key systems with the convenience of public key systems!

TLS uses KEM/DEM. TLS is used in https, so is extremely common.

Security Game for KEM

The game used to define security for KEM is very similar to that used to define security for key exchange, but we use the public key e and the ciphertext c in place of the transcript:



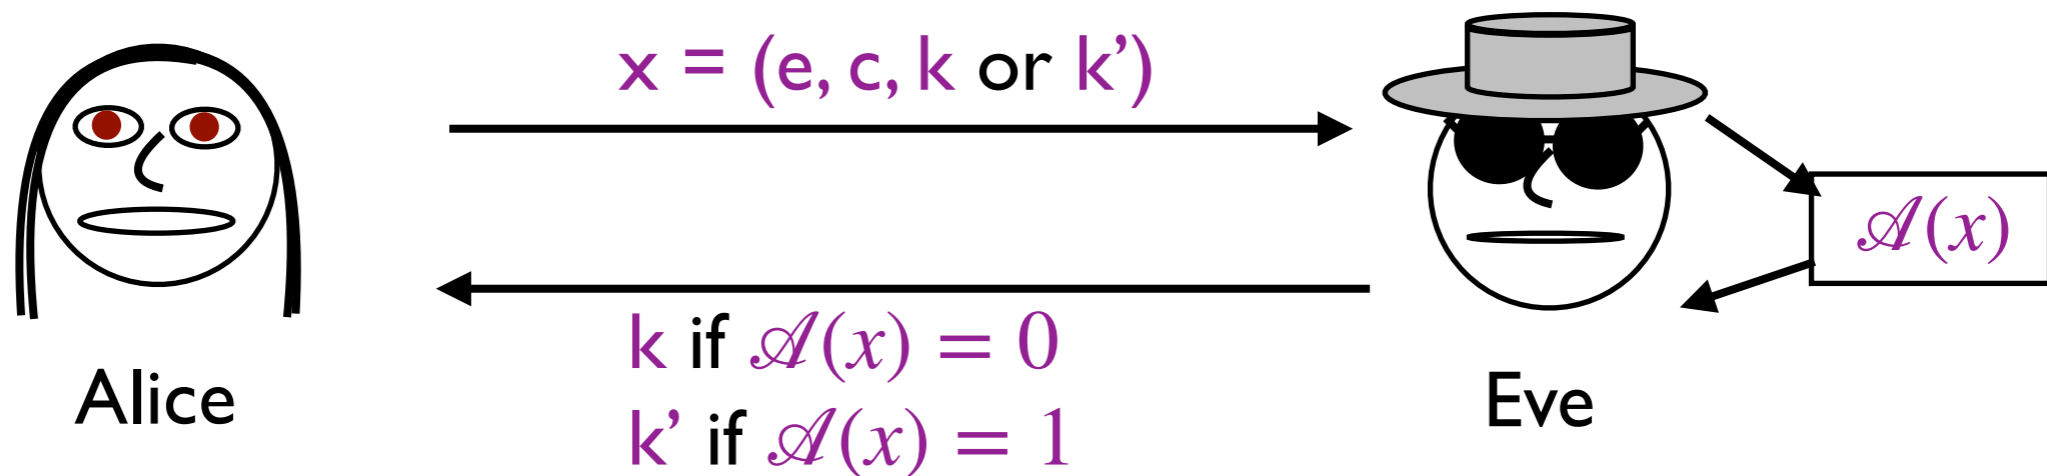
1. Alice runs **Encaps** to get c and k . She chooses a random bit r .
 - a. If $r = 0$, Alice sends Eve $x = (e, c, k)$.
 - b. If $r = 1$, Alice sends Eve $x = (e, c, k')$, where k' is a random bit string
2. Eve performs an attack $\mathcal{A}(x)$ to try to guess r (i.e., if Alice sent the real k or a random k'). She wins if she guesses right.

CPA Security for KEM

Definition: Consider a KEM (Gen , Encaps , Decaps). Suppose Gen produces public key e and $\text{Encaps}(e)$ produces ciphertext c and key k . Let k' be a uniformly randomly generated key. Then the KEM is **CPA secure** if for all attacks \mathcal{A} with a 1-bit output and taking as inputs e , c , and either k or k' ,

$$|\Pr(\mathcal{A}(e, c, k) = 1) - \Pr(\mathcal{A}(e, c, k') = 1)| \leq \epsilon(s)$$

with $\epsilon(s)$ negligible and the probabilities averaged over k' and over the randomness of \mathcal{A} , Gen , and Encaps .



Why is This CPA Security?

The term “CPA security” implies that it should be secure against chosen-plaintext attacks.

Why don't we need to give Eve an oracle to let her choose plaintexts?

Why is This CPA Security?

The term “CPA security” implies that it should be secure against chosen-plaintext attacks.

Why don't we need to give Eve an oracle to let her choose plaintexts?

Two reasons:

- Eve has access to e , so once again she can run **Encaps** if she wants.
- There is no plaintext here to choose. **Encaps** only takes e as input, no message.

Security of KEM/DEM

Theorem: If we have a KEM/DEM encryption scheme in which the KEM protocol is CPA secure and the DEM protocol is EAV secure, then the full KEM/DEM scheme is CPA secure.

Note: DEM only needs to be EAV secure.

Intuition and Proof Sketch:

The key from the CPA secure KEM looks the same to Eve as a random bit string. Therefore, if it used as the key in an EAV-secure encryption protocol, the protocol behaves the same way as if the key was fully random. That means that Eve cannot distinguish between messages.

Since Eve has access to the public key, this is all that is needed for CPA security of the public key protocol.

Note: Eve can create ciphertexts for the full protocol, but not for the DEM with the specific key used by Alice.

