

CMSC/Math 456: Cryptography (Fall 2023)

Lecture 23

Daniel Gottesman

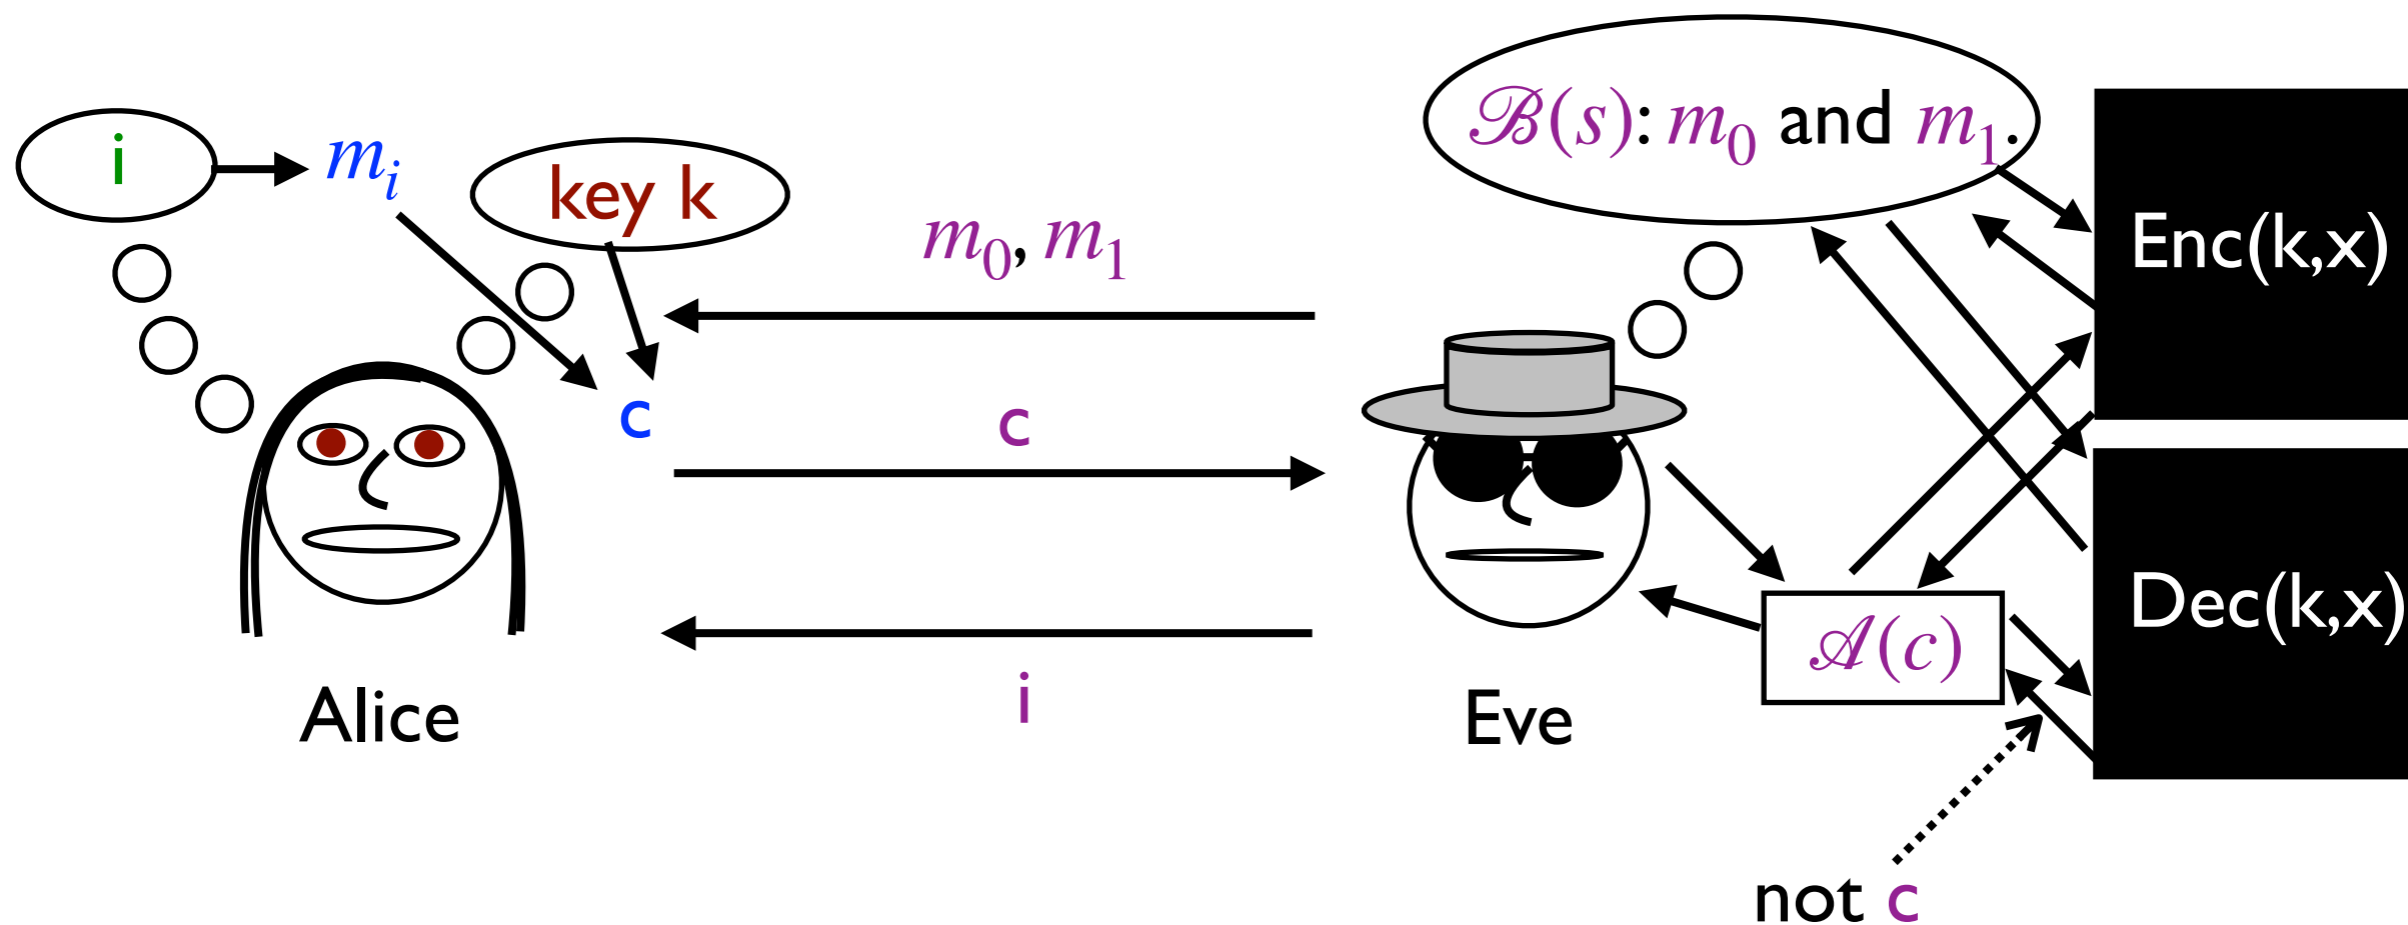
Administrative

Problem set #9 is out but you will have *two weeks* to do it. (Due Nov. 30). The solution set for Problem Set #7 is also available on ELMS.

Reminder: No class next week.

CCA Security

A protocol is CCA-secure if Eve cannot guess the encrypted message even with access to both encryption and decryption oracles.




Unforgeability

Definition: An encryption protocol $(\text{Gen}, \text{Enc}, \text{Dec})$ with security parameter s is **unforgeable** if, for any polynomial-time attack \mathcal{A} with oracle access to $\text{Enc}(k, m)$, where \mathcal{A} outputs \hat{c} with $\hat{m} = \text{Dec}(\hat{c})$ such that \mathcal{A} never queried the oracle for $m = \hat{m}$,

$$\Pr(\hat{c} \text{ is valid}) \leq \epsilon(s)$$

where $\epsilon(s)$ is a negligible function and the probability is averaged over k generated by Gen and the randomness used in any of the functions.



Authenticated Encryption

Definition: A private-key encryption scheme provides **authenticated encryption** if it is CCA-secure and unforgeable.

Theorem: If (Enc, Dec) is a CPA-secure encryption scheme and $(\text{Mac}, \text{Vrfy})$ is a strongly secure MAC, then the following encryption scheme is an authenticated encryption protocol:

Enc': Given message m and keys k and k' , the ciphertext is $(\text{Enc}(k, m), \text{Mac}(k', \text{Enc}(k, m)))$.

Dec': Given ciphertext (c, t) and keys (k, k') , output **Invalid** if $\text{Vrfy}(k', c, t)$ is invalid. Otherwise decrypt c to $m = \text{Dec}(k, c)$ and output m .

Security of the Protocol

Why does the encrypt-then-authenticate strategy produce a secure authenticated encryption scheme?

Security of the Protocol

Why does the encrypt-then-authenticate strategy produce a secure authenticated encryption scheme?

- The protocol is certainly unforgeable because any new (not produced by the **Enc** oracle) ciphertext an adversary might produce will be invalid for the MAC (since the MAC is strongly secure)

Security of the Protocol

Why does the encrypt-then-authenticate strategy produce a secure authenticated encryption scheme?

- The protocol is certainly unforgeable because any new (not produced by the **Enc** oracle) ciphertext an adversary might produce will be invalid for the MAC (since the MAC is strongly secure)
- Consider an attack $(\mathcal{A}, \mathcal{B})$. It may make queries to the decryption oracle, but all of them (except with negligible probability or any ciphertexts produced by **Enc**) will be invalid.

Security of the Protocol

Why does the encrypt-then-authenticate strategy produce a secure authenticated encryption scheme?

- The protocol is certainly unforgeable because any new (not produced by the **Enc** oracle) ciphertext an adversary might produce will be invalid for the MAC (since the MAC is strongly secure)
- Consider an attack $(\mathcal{A}, \mathcal{B})$. It may make queries to the decryption oracle, but all of them (except with negligible probability or any ciphertexts produced by **Enc**) will be invalid.
- Therefore any **Dec** queries are useless to \mathcal{A} and can be easily simulated by Eve herself.

Security of the Protocol

Why does the encrypt-then-authenticate strategy produce a secure authenticated encryption scheme?

- The protocol is certainly unforgeable because any new (not produced by the **Enc** oracle) ciphertext an adversary might produce will be invalid for the MAC (since the MAC is strongly secure)
- Consider an attack $(\mathcal{A}, \mathcal{B})$. It may make queries to the decryption oracle, but all of them (except with negligible probability or any ciphertexts produced by **Enc**) will be invalid.
- Therefore any **Dec** queries are useless to \mathcal{A} and can be easily simulated by Eve herself.
- Thus, the attack $(\mathcal{A}, \mathcal{B})$ could be performed by a CPA adversary with access only to the **Enc** oracle.

Security of the Protocol

Why does the encrypt-then-authenticate strategy produce a secure authenticated encryption scheme?

- The protocol is certainly unforgeable because any new (not produced by the **Enc** oracle) ciphertext an adversary might produce will be invalid for the MAC (since the MAC is strongly secure)
- Consider an attack $(\mathcal{A}, \mathcal{B})$. It may make queries to the decryption oracle, but all of them (except with negligible probability or any ciphertexts produced by **Enc**) will be invalid.
- Therefore any **Dec** queries are useless to \mathcal{A} and can be easily simulated by Eve herself.
- Thus, the attack $(\mathcal{A}, \mathcal{B})$ could be performed by a CPA adversary with access only to the **Enc** oracle.
- Since the encryption protocol is CPA secure, Eve cannot distinguish messages except with negligible probability.

CCA Security for Public Key Protocols

Since message authentication is a private-key protocol, we can't combine it with a public-key protocol straightforwardly. In particular, the notion of authenticated encryption needs some modification. To understand what should replace it, we will need to discuss **digital signatures**, the next topic.

But for now, let us focus only on CCA security of public key protocols.

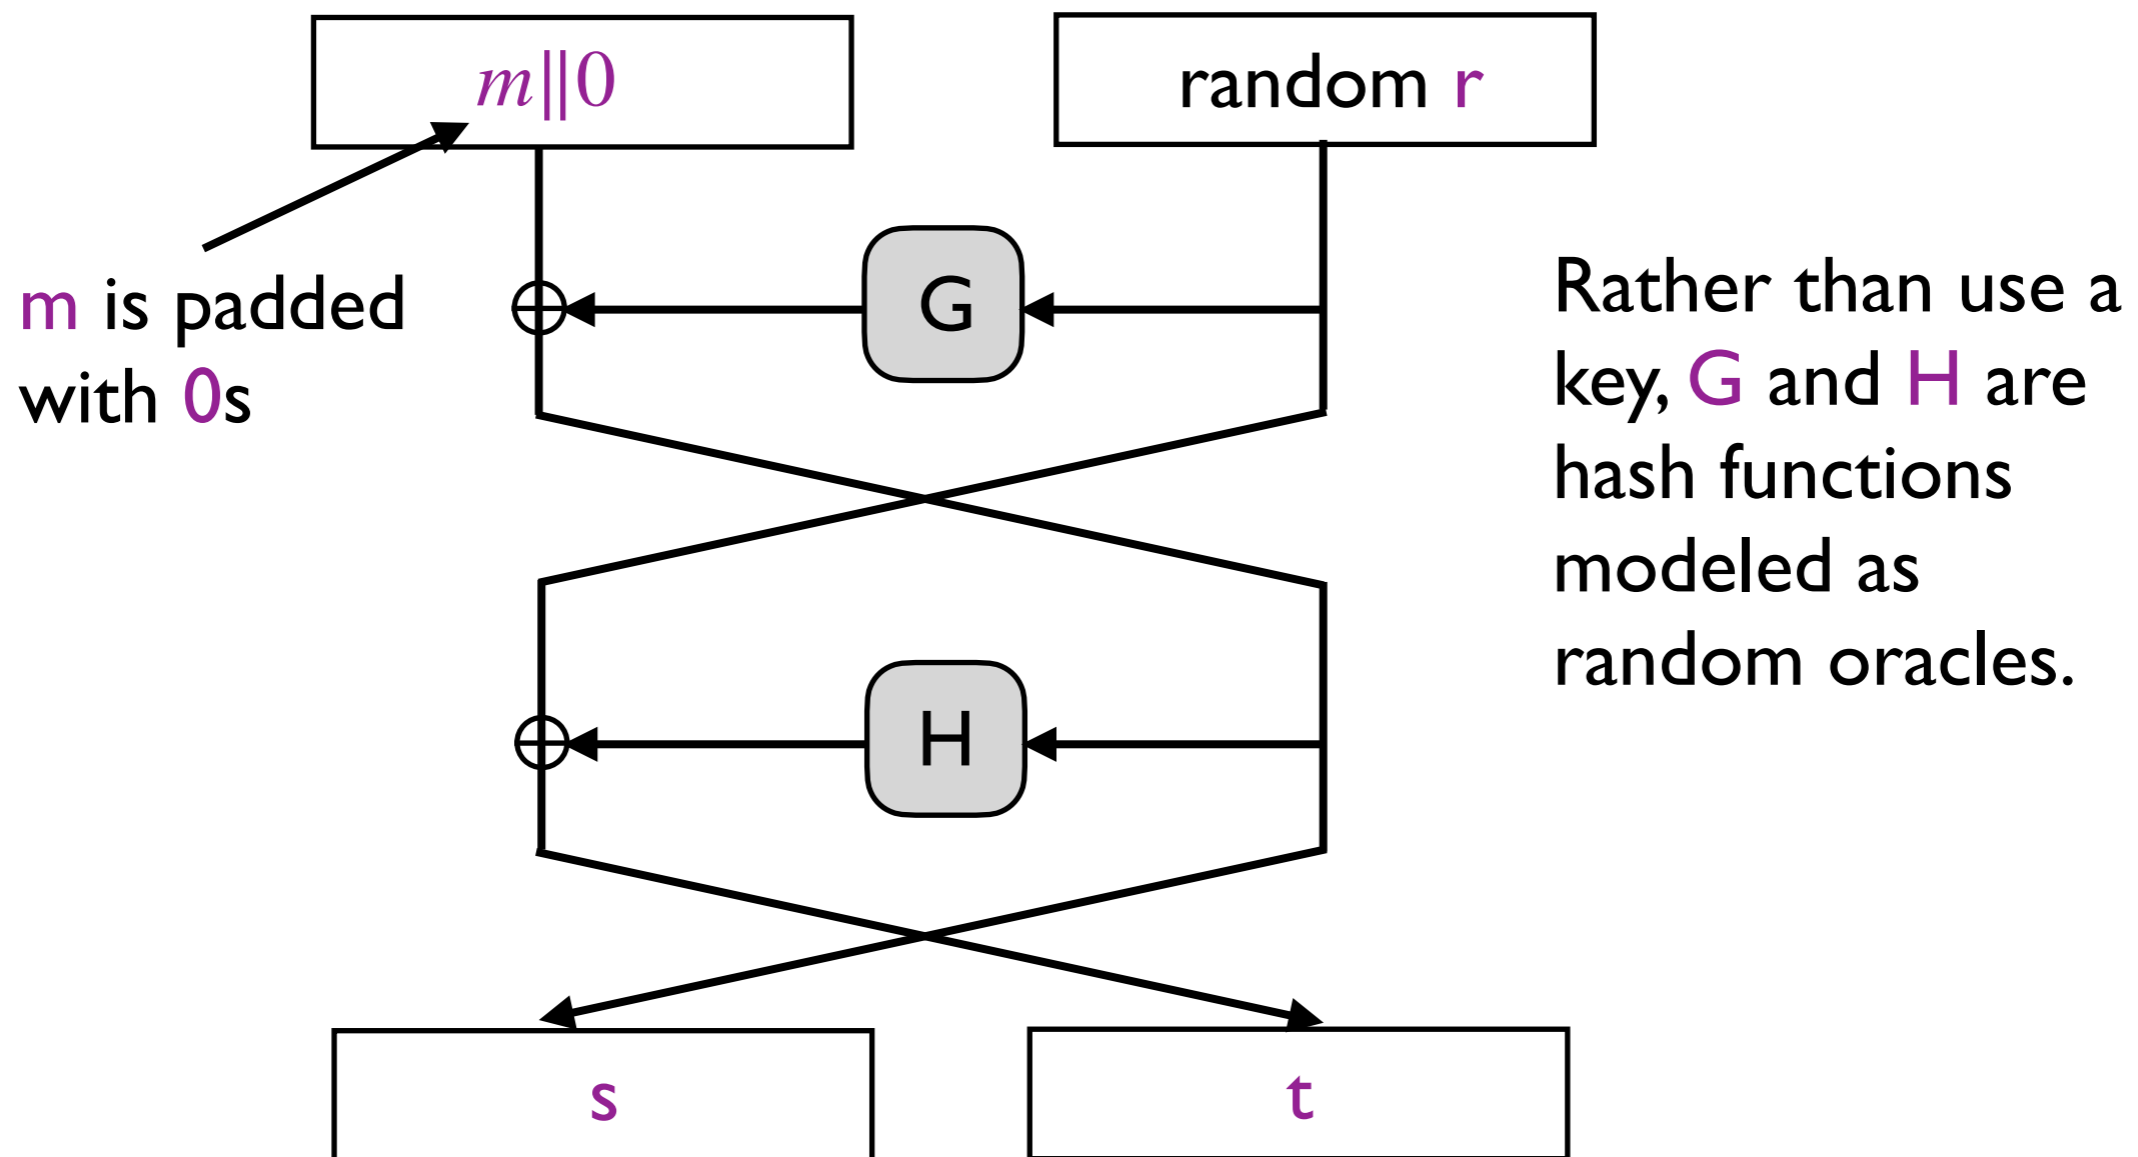
RSA and El Gamal are **malleable**, so not CCA-secure.

There are two basic approaches we can take:

- Modify CPA-secure public key encryption protocols to make them CCA-secure.
- Use KEM/DEM with a CCA-secure private key scheme.

RSA-OAEP


In RSA-OAEP (“Optimal asymmetric encryption padding”), the message m is put through a Feistel network.



Then $s||t$ is encrypted with RSA: ciphertext $(s||t)^e \bmod N$.

RSA-OAEP Decryption

Decryption recovers $(s||t) = c^d \bmod N$ and then runs the Feistel network backwards.



Bob rejects the ciphertext as invalid unless the correct number of padding 0s are present.

RSA-OAEP

RSA-OAEP is CCA secure if G and H are modeled as random oracles and the RSA assumption is true (It is hard to find x such that $x^e = y \pmod N$ for random y).

Because G and H are considered as random oracles, Eve has to know nearly every bit of r and t in order to be able to query them properly. The RSA assumption means that Eve is missing some of this information.

As with the authenticated encryption case, Eve cannot come up with a ciphertext that gives the correct padding.

RSA-OAEP

RSA-OAEP is CCA secure if G and H are modeled as random oracles and the RSA assumption is true (It is hard to find x such that $x^e = y \pmod N$ for random y).

Because G and H are considered as random oracles, Eve has to know nearly every bit of r and t in order to be able to query them properly. The RSA assumption means that Eve is missing some of this information.

As with the authenticated encryption case, Eve cannot come up with a ciphertext that gives the correct padding.

Caveat: There are two possible “invalid” responses: 1) padding is wrong and 2) the message is not in the right form to be $(s||t)$ (i.e., too large for the total length of s and t). These two types of errors must be indistinguishable (including via side-channel information) or the scheme can be broken.

KEM Approach

When we combine a CCA-secure KEM with a CCA-secure DEM/private key system, we get a **CCA-secure public key system**.

Reminder:

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

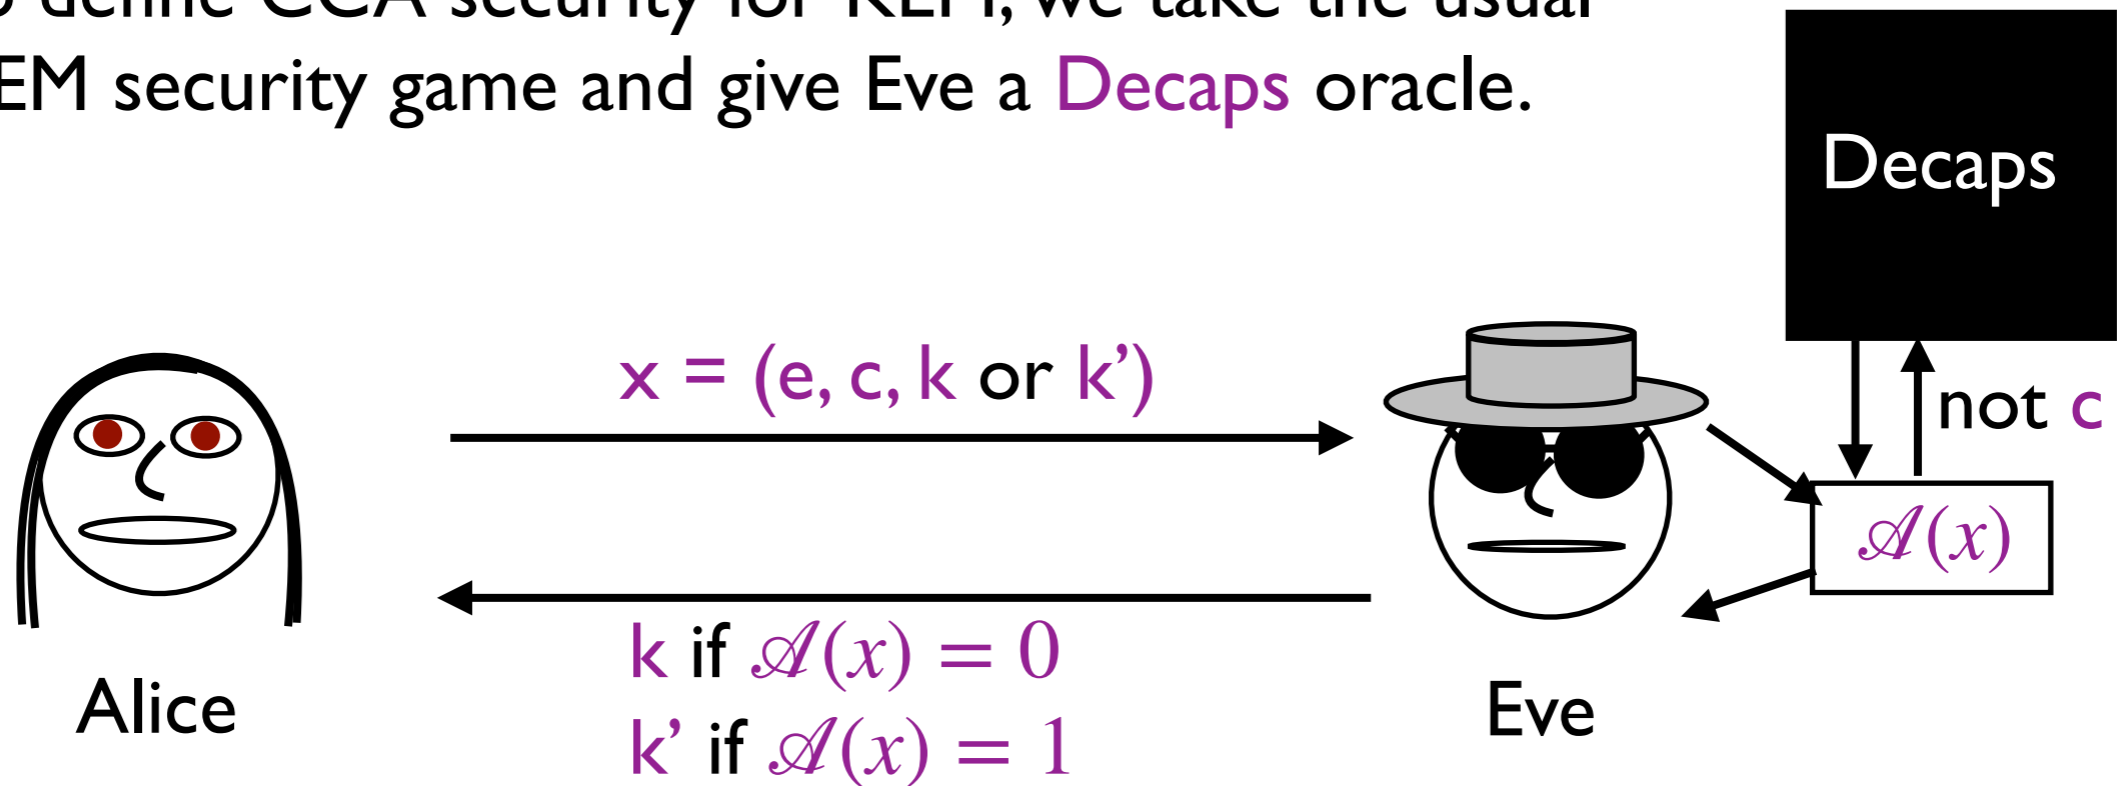
Gen is the **key generation algorithm**. It takes as input **s**, the **security parameter**, and outputs a **public key, private key pair** $(e, d) \in \{0,1\}^* \times \{0,1\}^*$.

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 $k' \in \{0,1\}^{\ell(s)}$.

CCA Security Game for KEM

To define CCA security for KEM, we take the usual KEM security game and give Eve a **Decaps** oracle.



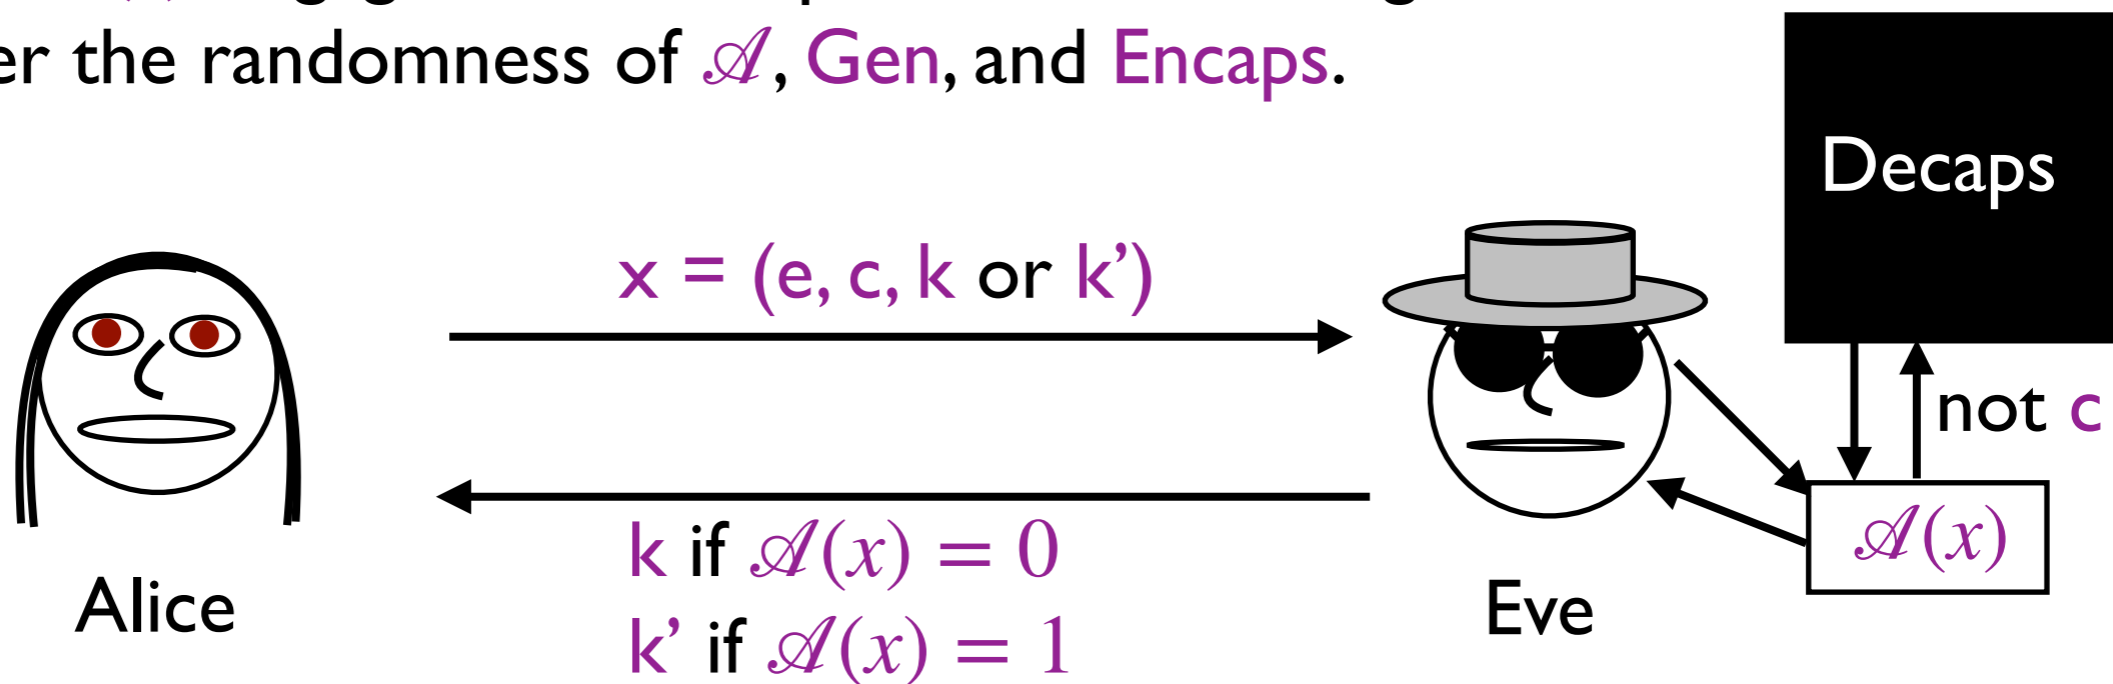
1. Alice runs **Encaps** to get c and k . She chooses a random bit r .
 - a. If $r = 0$, Alice sends Eve $x = (e, c, k)$.
 - b. If $r = 1$, Alice sends Eve $x = (e, c, k')$, where k' is a random bit string
2. Eve performs an attack $\mathcal{A}(x)$ to try to guess r (i.e., if Alice sent the real k or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

CCA 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 **CCA secure** if for all attacks \mathcal{A} taking as inputs e , c , and either k or k' , and with oracle access to Decaps except for c ,

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

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



CCA-Secure KEM/DEM

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

Note: Both protocols need to be **CCA-secure**. (Whereas for CPA-security, the DEM only needed to be EAV-secure.)

The proof is similar to the CPA-security case:


- Because the KEM is CCA-secure, the key generated by it looks like a random bit string to Eve, even with a **Dec** oracle.
- With a random key, the DEM is CCA-secure, so Eve can't distinguish between messages.

DHIES


The usual Diffie-Hellman-based KEM is **CCA-secure** assuming:

- A stronger version of the Diffie-Hellman security assumption.
- $H(x)$ is modeled as a random oracle.

When combined with a CCA-secure private key system, this is **DHIES (Diffie-Hellman Integrated Encryption System)**.



Alice




Bob

DHIES

The usual Diffie-Hellman-based KEM is **CCA-secure** assuming:

- A stronger version of the Diffie-Hellman security assumption.
- $H(x)$ is modeled as a random oracle.

When combined with a CCA-secure private key system, this is **DHIES (Diffie-Hellman Integrated Encryption System)**.

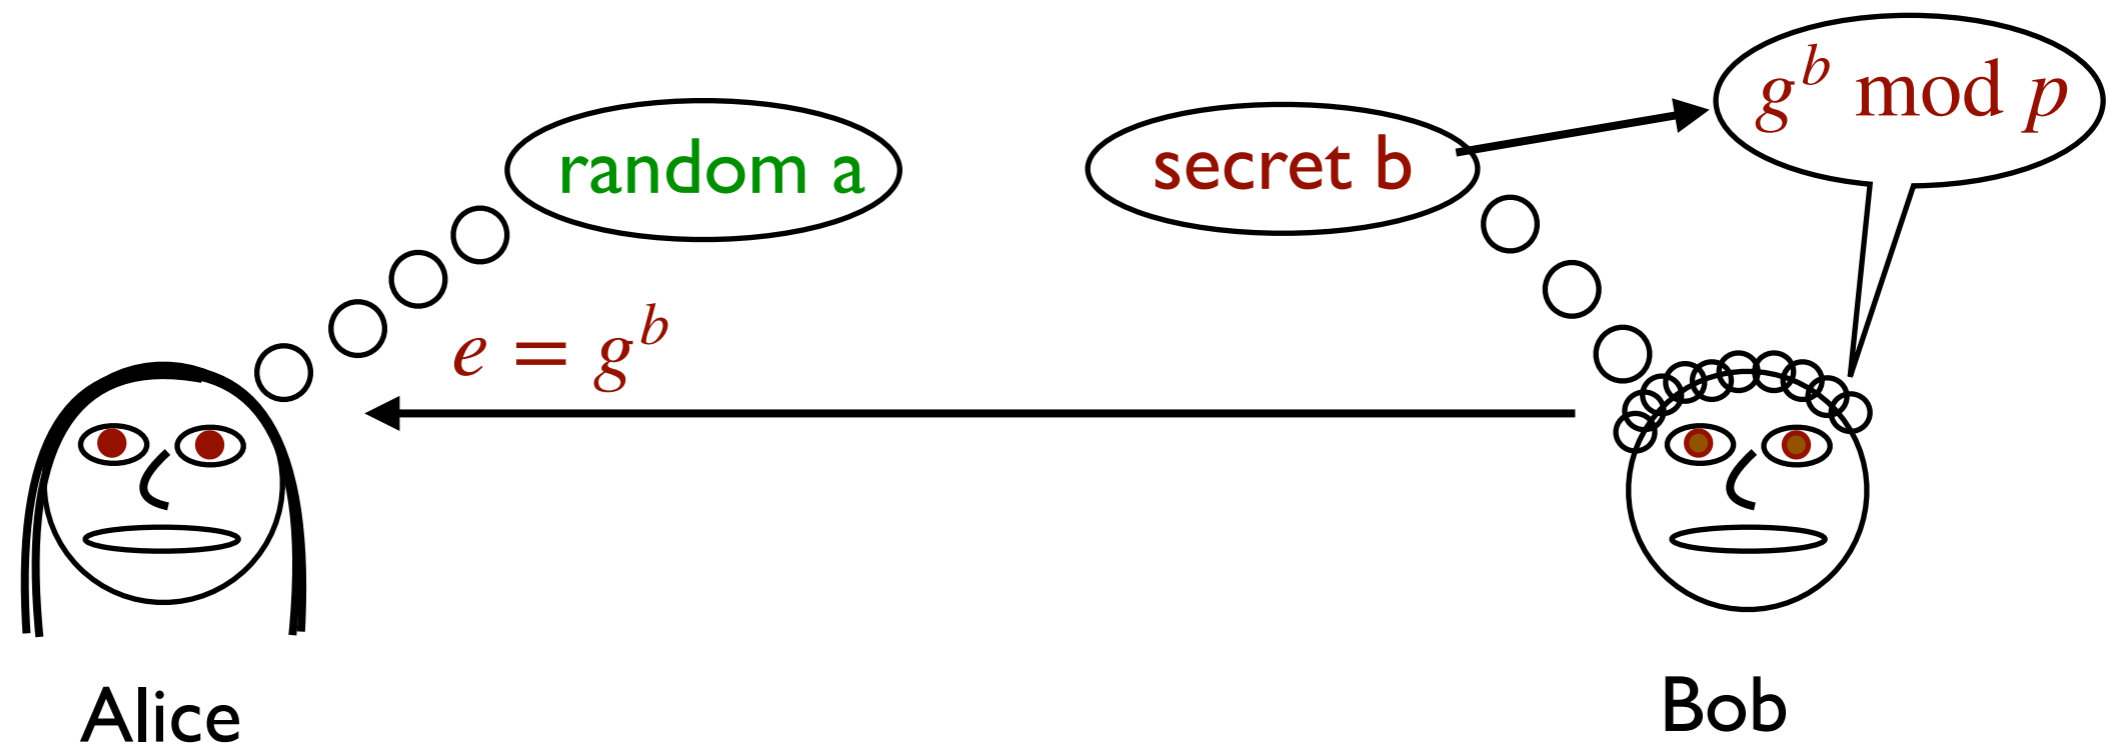


DHIES

The usual Diffie-Hellman-based KEM is **CCA-secure** assuming:

- A stronger version of the Diffie-Hellman security assumption.
- $H(x)$ is modeled as a random oracle.

When combined with a CCA-secure private key system, this is **DHIES (Diffie-Hellman Integrated Encryption System)**.

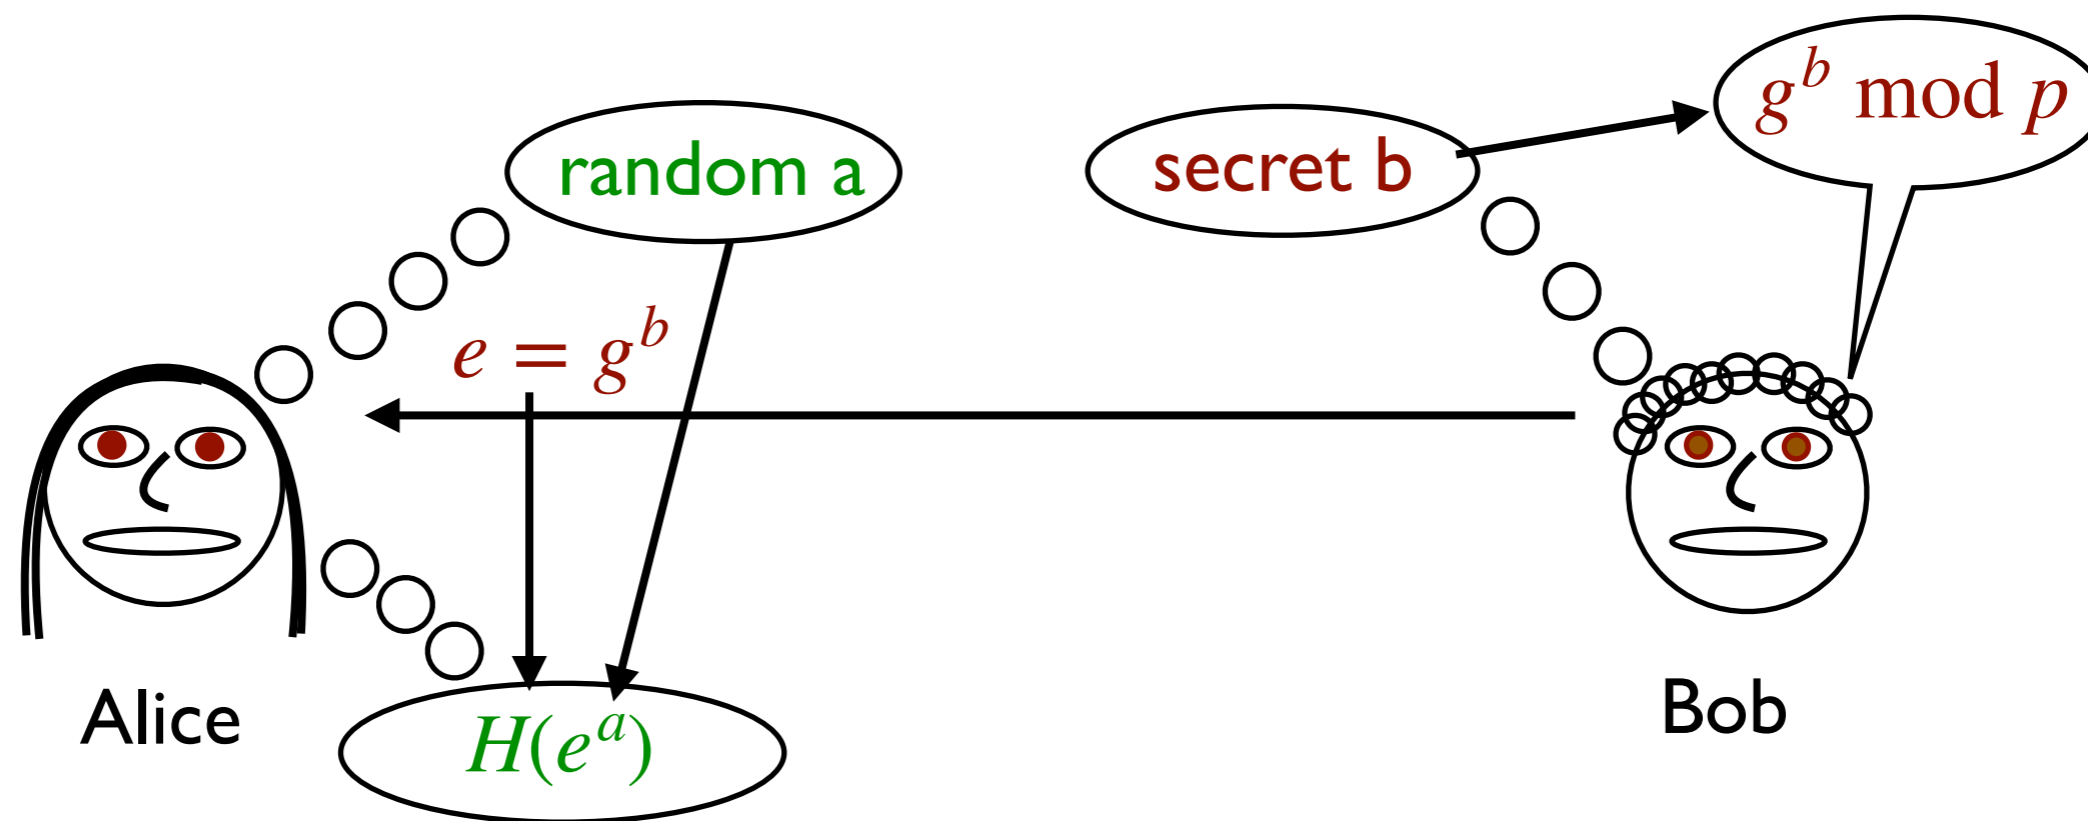


DHIES

The usual Diffie-Hellman-based KEM is **CCA-secure** assuming:

- A stronger version of the Diffie-Hellman security assumption.
- $H(x)$ is modeled as a random oracle.

When combined with a CCA-secure private key system, this is **DHIES (Diffie-Hellman Integrated Encryption System)**.

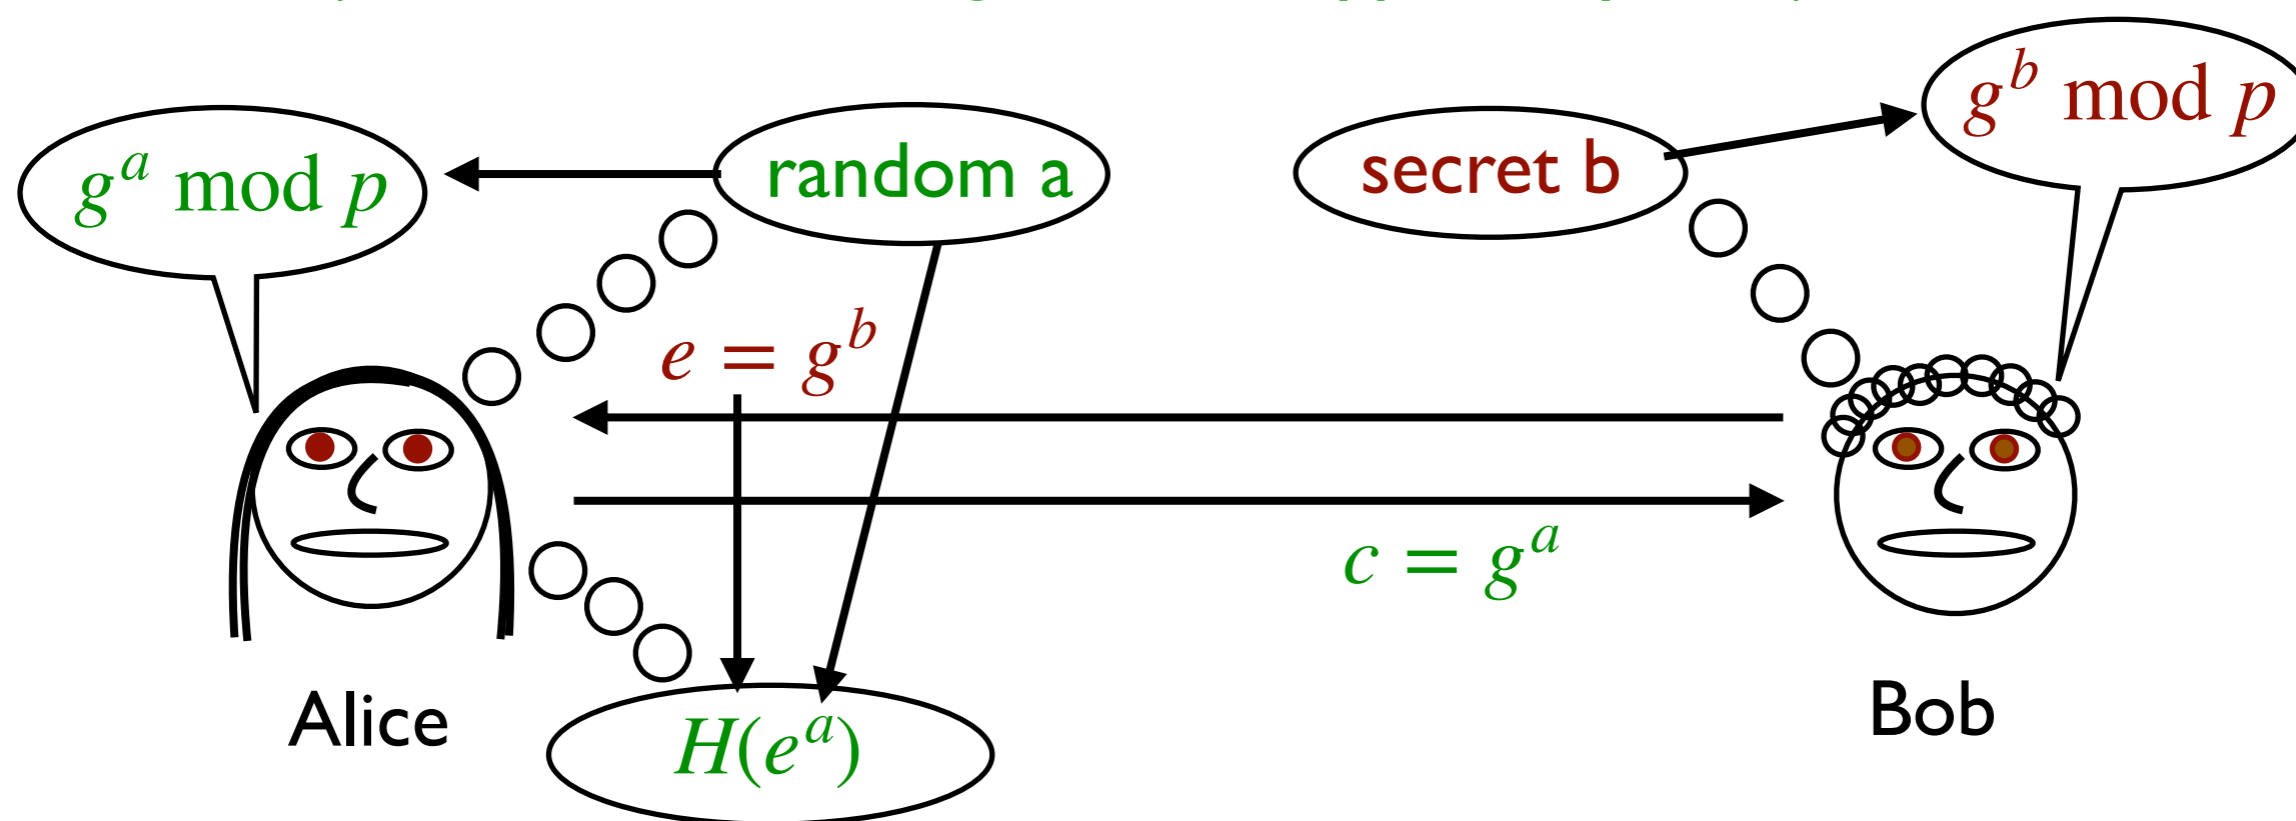


DHIES

The usual Diffie-Hellman-based KEM is **CCA-secure** assuming:

- A stronger version of the Diffie-Hellman security assumption.
- $H(x)$ is modeled as a random oracle.

When combined with a CCA-secure private key system, this is **DHIES (Diffie-Hellman Integrated Encryption System)**.

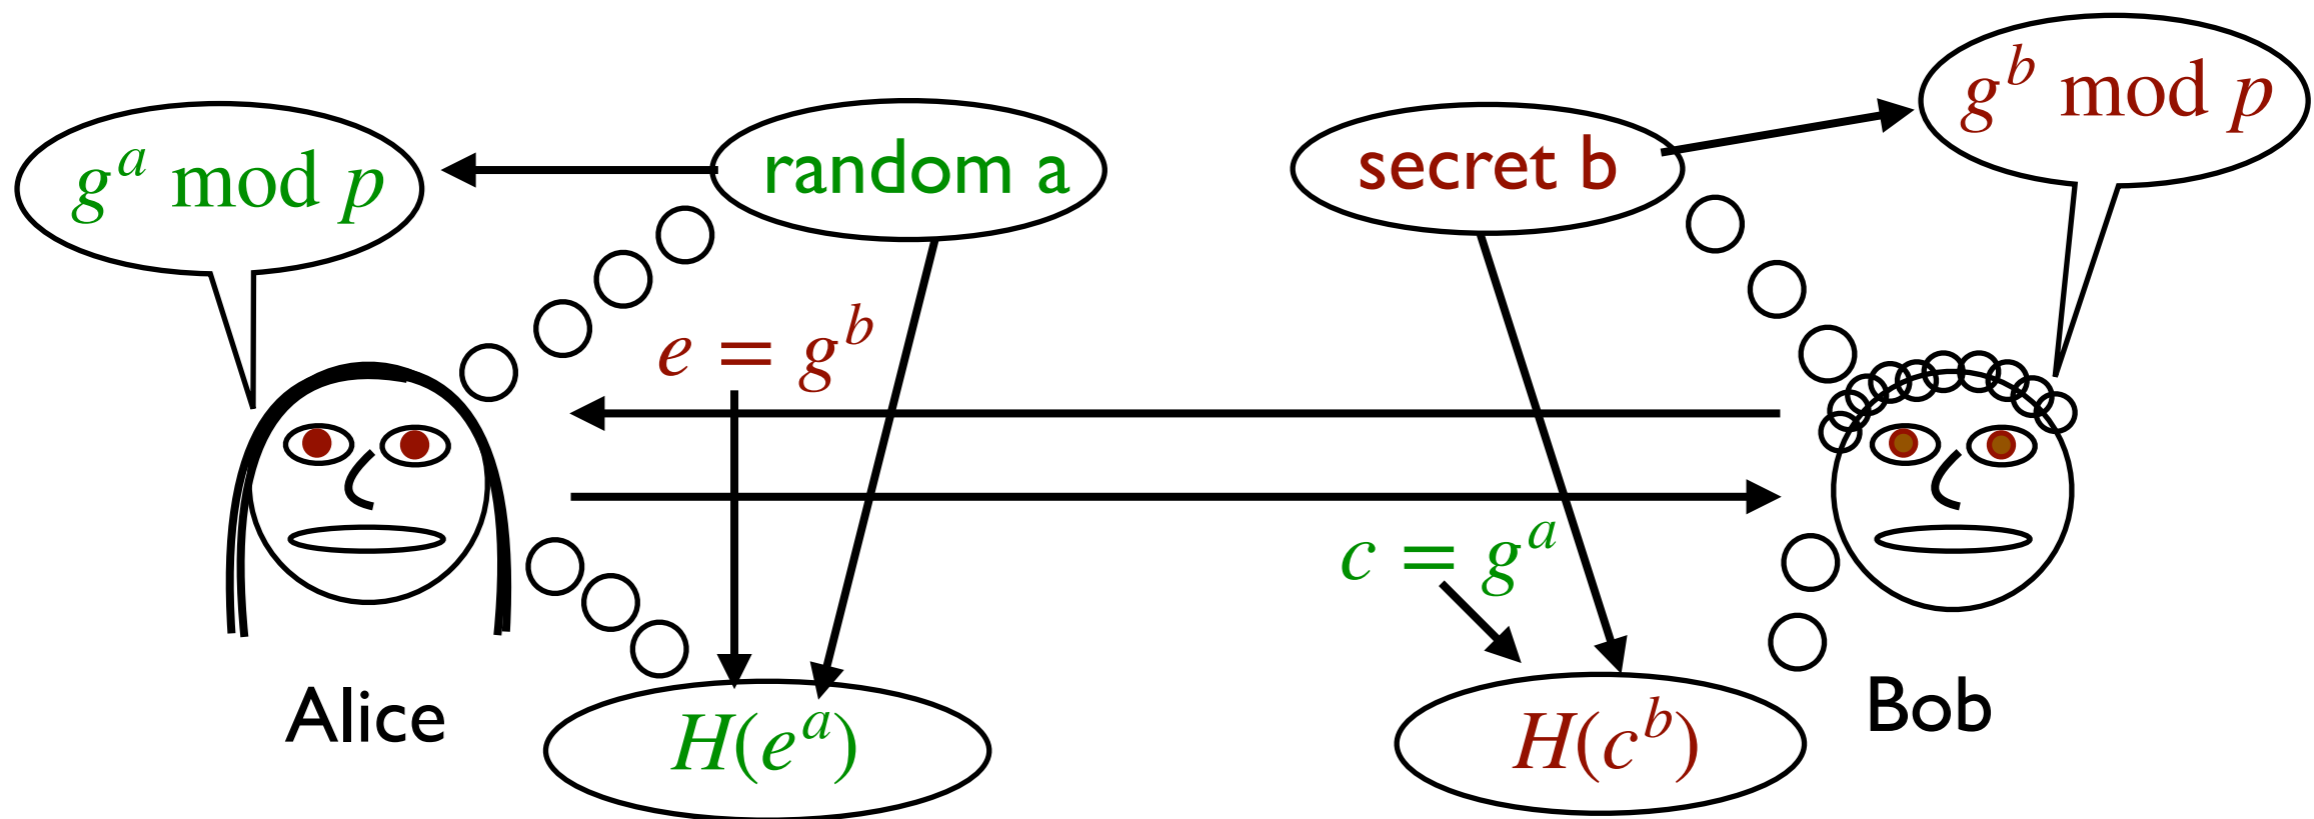


DHIES

The usual Diffie-Hellman-based KEM is **CCA-secure** assuming:

- A stronger version of the Diffie-Hellman security assumption.
- $H(x)$ is modeled as a random oracle.

When combined with a CCA-secure private key system, this is **DHIES (Diffie-Hellman Integrated Encryption System)**.

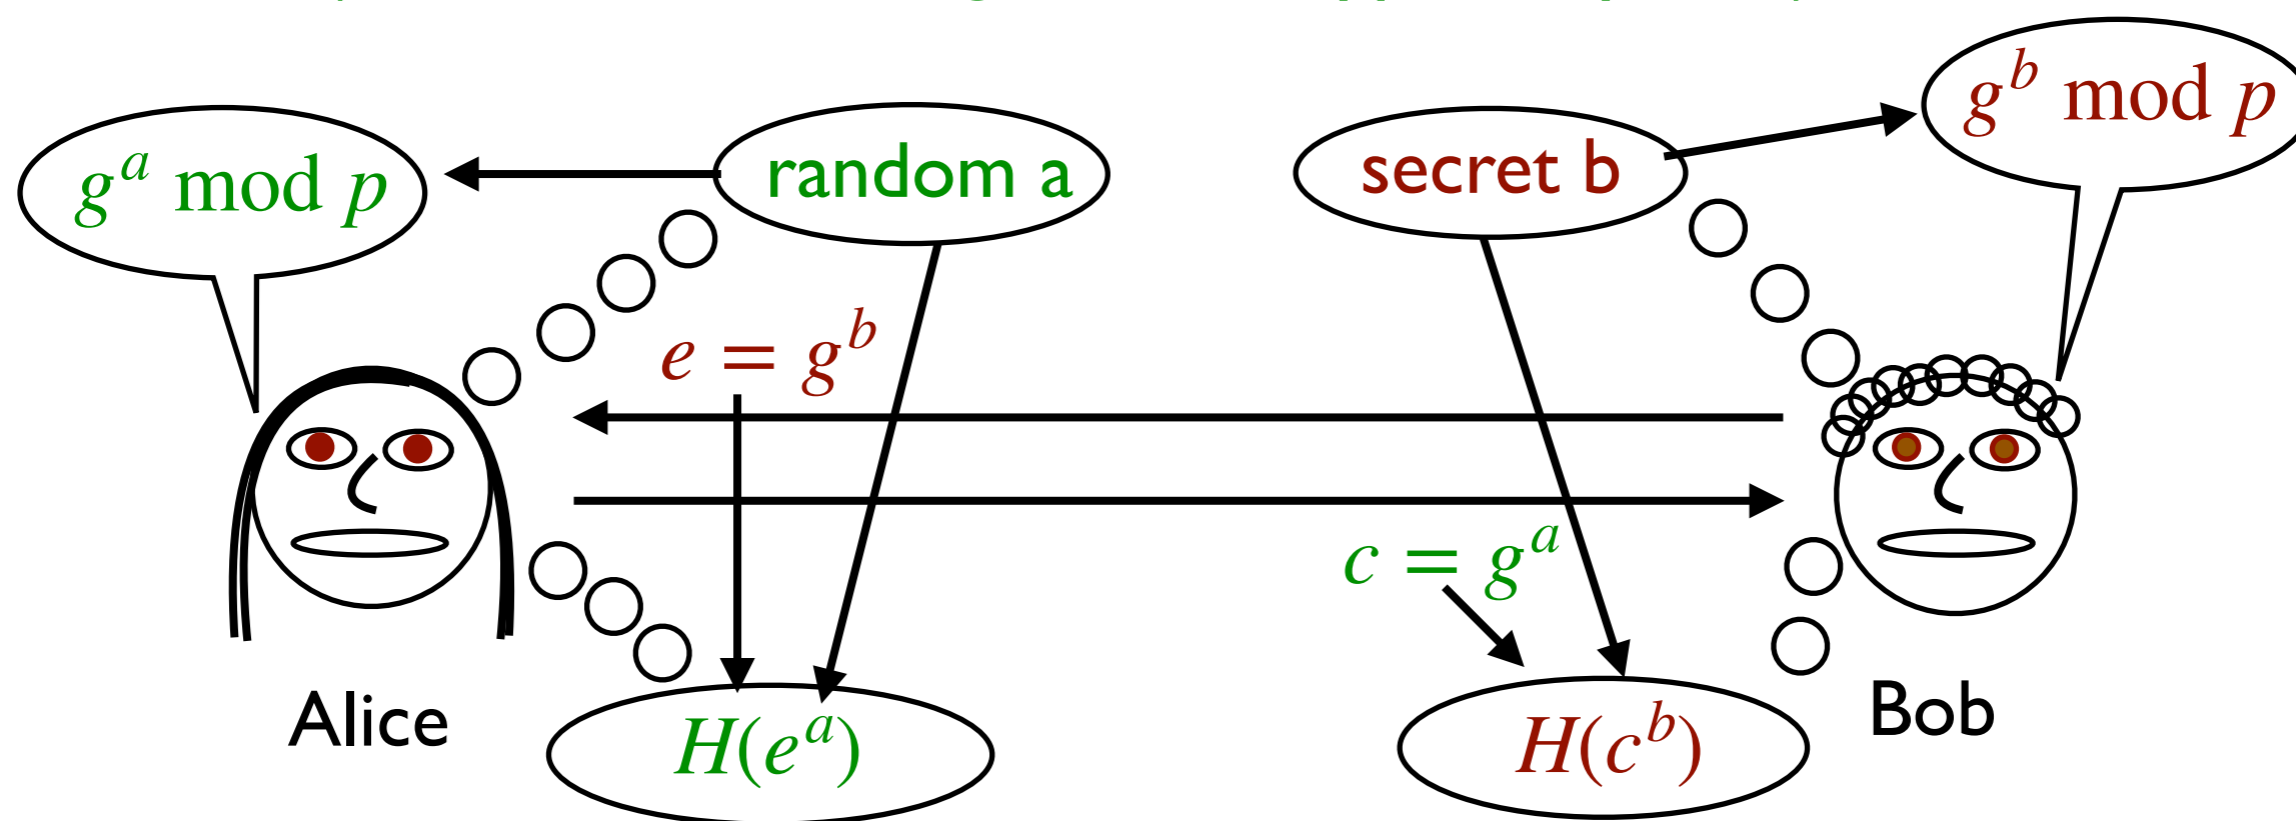


DHIES

The usual Diffie-Hellman-based KEM is **CCA-secure** assuming:

- A stronger version of the Diffie-Hellman security assumption.
- $H(x)$ is modeled as a random oracle.

When combined with a CCA-secure private key system, this is **DHIES (Diffie-Hellman Integrated Encryption System)**.



One change: Check c is a possible power of g (i.e., its order is r , not $2r = p-1$).

RSA-Based CCA-Secure KEM

Gen: Pick a (public) **key derivation function** $H(x)$, then as usual for RSA, i.e., generate two random primes p and q which are s bits long. Let $N = pq$. Choose $e, d \in \mathbb{Z}_N^*$ s.t. $ed = 1 \pmod{\varphi(N)}$. The **public key** is (N, e) and the **private key** is (N, d) .

Encaps: Choose random x . The **ciphertext** is $c = x^e \pmod N$ and the **key** is $H(x)$.

Decaps: Given c and d , compute $x' = c^d \pmod N$. Then the key is $H(x')$.

This is again the standard RSA-based KEM. It is **CCA-secure** assuming:

- The standard RSA security assumption.
- $H(x)$ is modeled as a random oracle.

Why Are These KEMs CCA Secure?

In both cases, the underlying public key scheme (El Gamal or RSA) is malleable.

Why Are These KEMs CCA Secure?

In both cases, the underlying public key scheme (El Gamal or RSA) is malleable.

However, when H acts like a random oracle, it breaks the relationship between ciphertexts and encapsulated keys.

Why Are These KEMs CCA Secure?

In both cases, the underlying public key scheme (El Gamal or RSA) is malleable.

However, when H acts like a random oracle, it breaks the relationship between ciphertexts and encapsulated keys.

In particular, while Eve can easily come up with other valid ciphertexts, they don't tell her anything about the actual key since the output of $H(x)$ is uncorrelated with $H(x')$ when $x \neq x'$.

Why Are These KEMs CCA Secure?

In both cases, the underlying public key scheme (El Gamal or RSA) is malleable.

However, when H acts like a random oracle, it breaks the relationship between ciphertexts and encapsulated keys.

In particular, while Eve can easily come up with other valid ciphertexts, they don't tell her anything about the actual key since the output of $H(x)$ is uncorrelated with $H(x')$ when $x \neq x'$.

For example, suppose Eve sees the ciphertext $c = x^e \bmod N$ for RSA-KEM and submits $c' = 2^e c = (2x)^e \bmod N$ to the **Decaps** oracle. She then gets $H(2x)$ which is unrelated to $H(x)$.

Why Are These KEMs CCA Secure?

In both cases, the underlying public key scheme (El Gamal or RSA) is malleable.

However, when H acts like a random oracle, it breaks the relationship between ciphertexts and encapsulated keys.


In particular, while Eve can easily come up with other valid ciphertexts, they don't tell her anything about the actual key since the output of $H(x)$ is uncorrelated with $H(x')$ when $x \neq x'$.

For example, suppose Eve sees the ciphertext $c = x^e \bmod N$ for RSA-KEM and submits $c' = 2^e c = (2x)^e \bmod N$ to the **Decaps** oracle. She then gets $H(2x)$ which is unrelated to $H(x)$.

This is a **different mechanism** to get CCA security than we saw before.

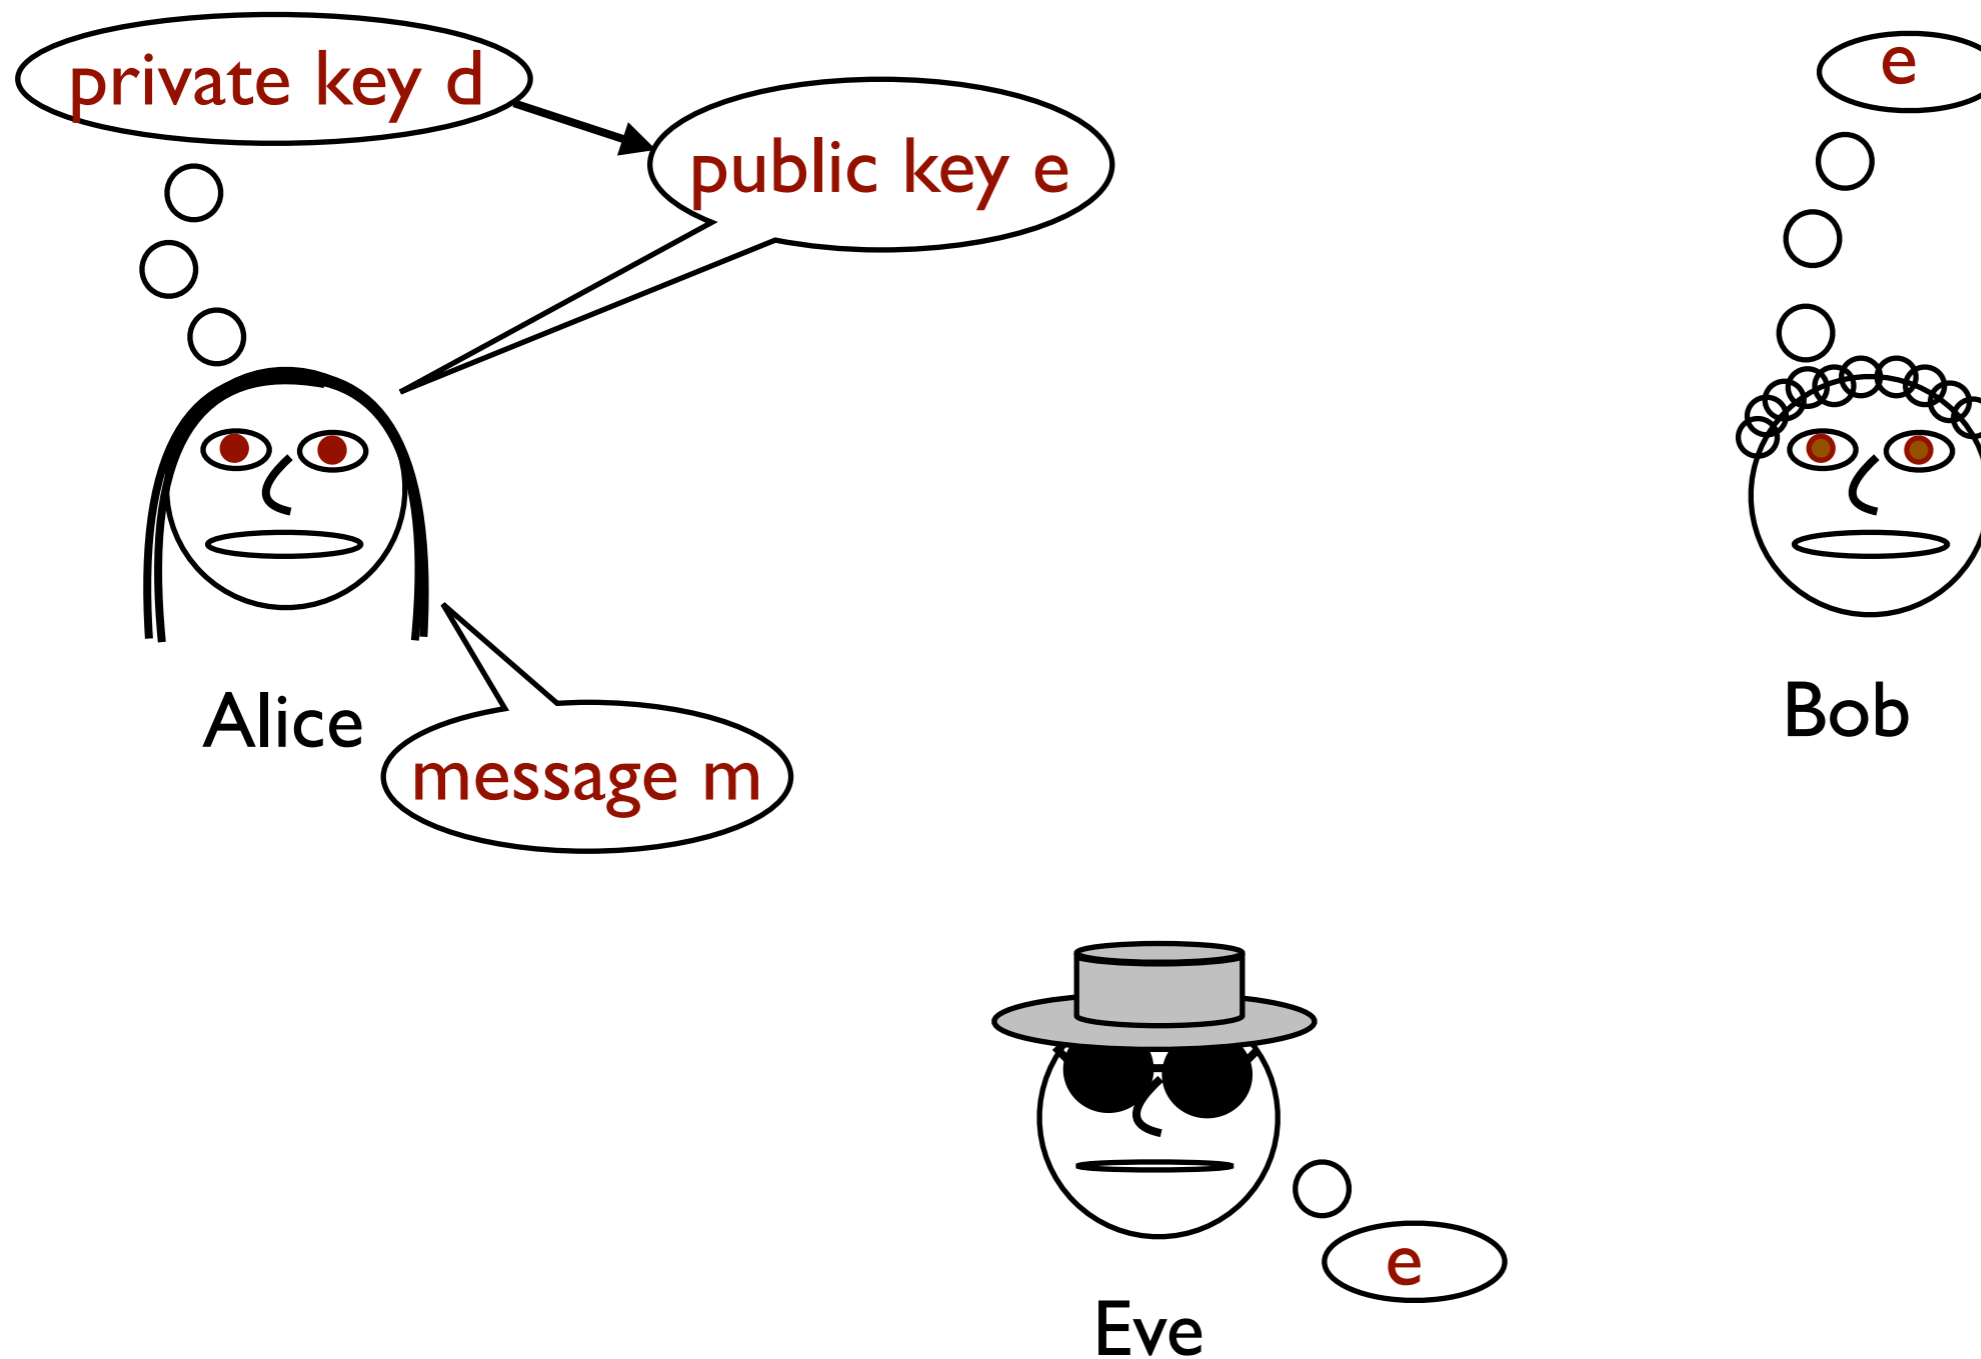
Digital Signatures

Digital signatures are a public key version of MACs.



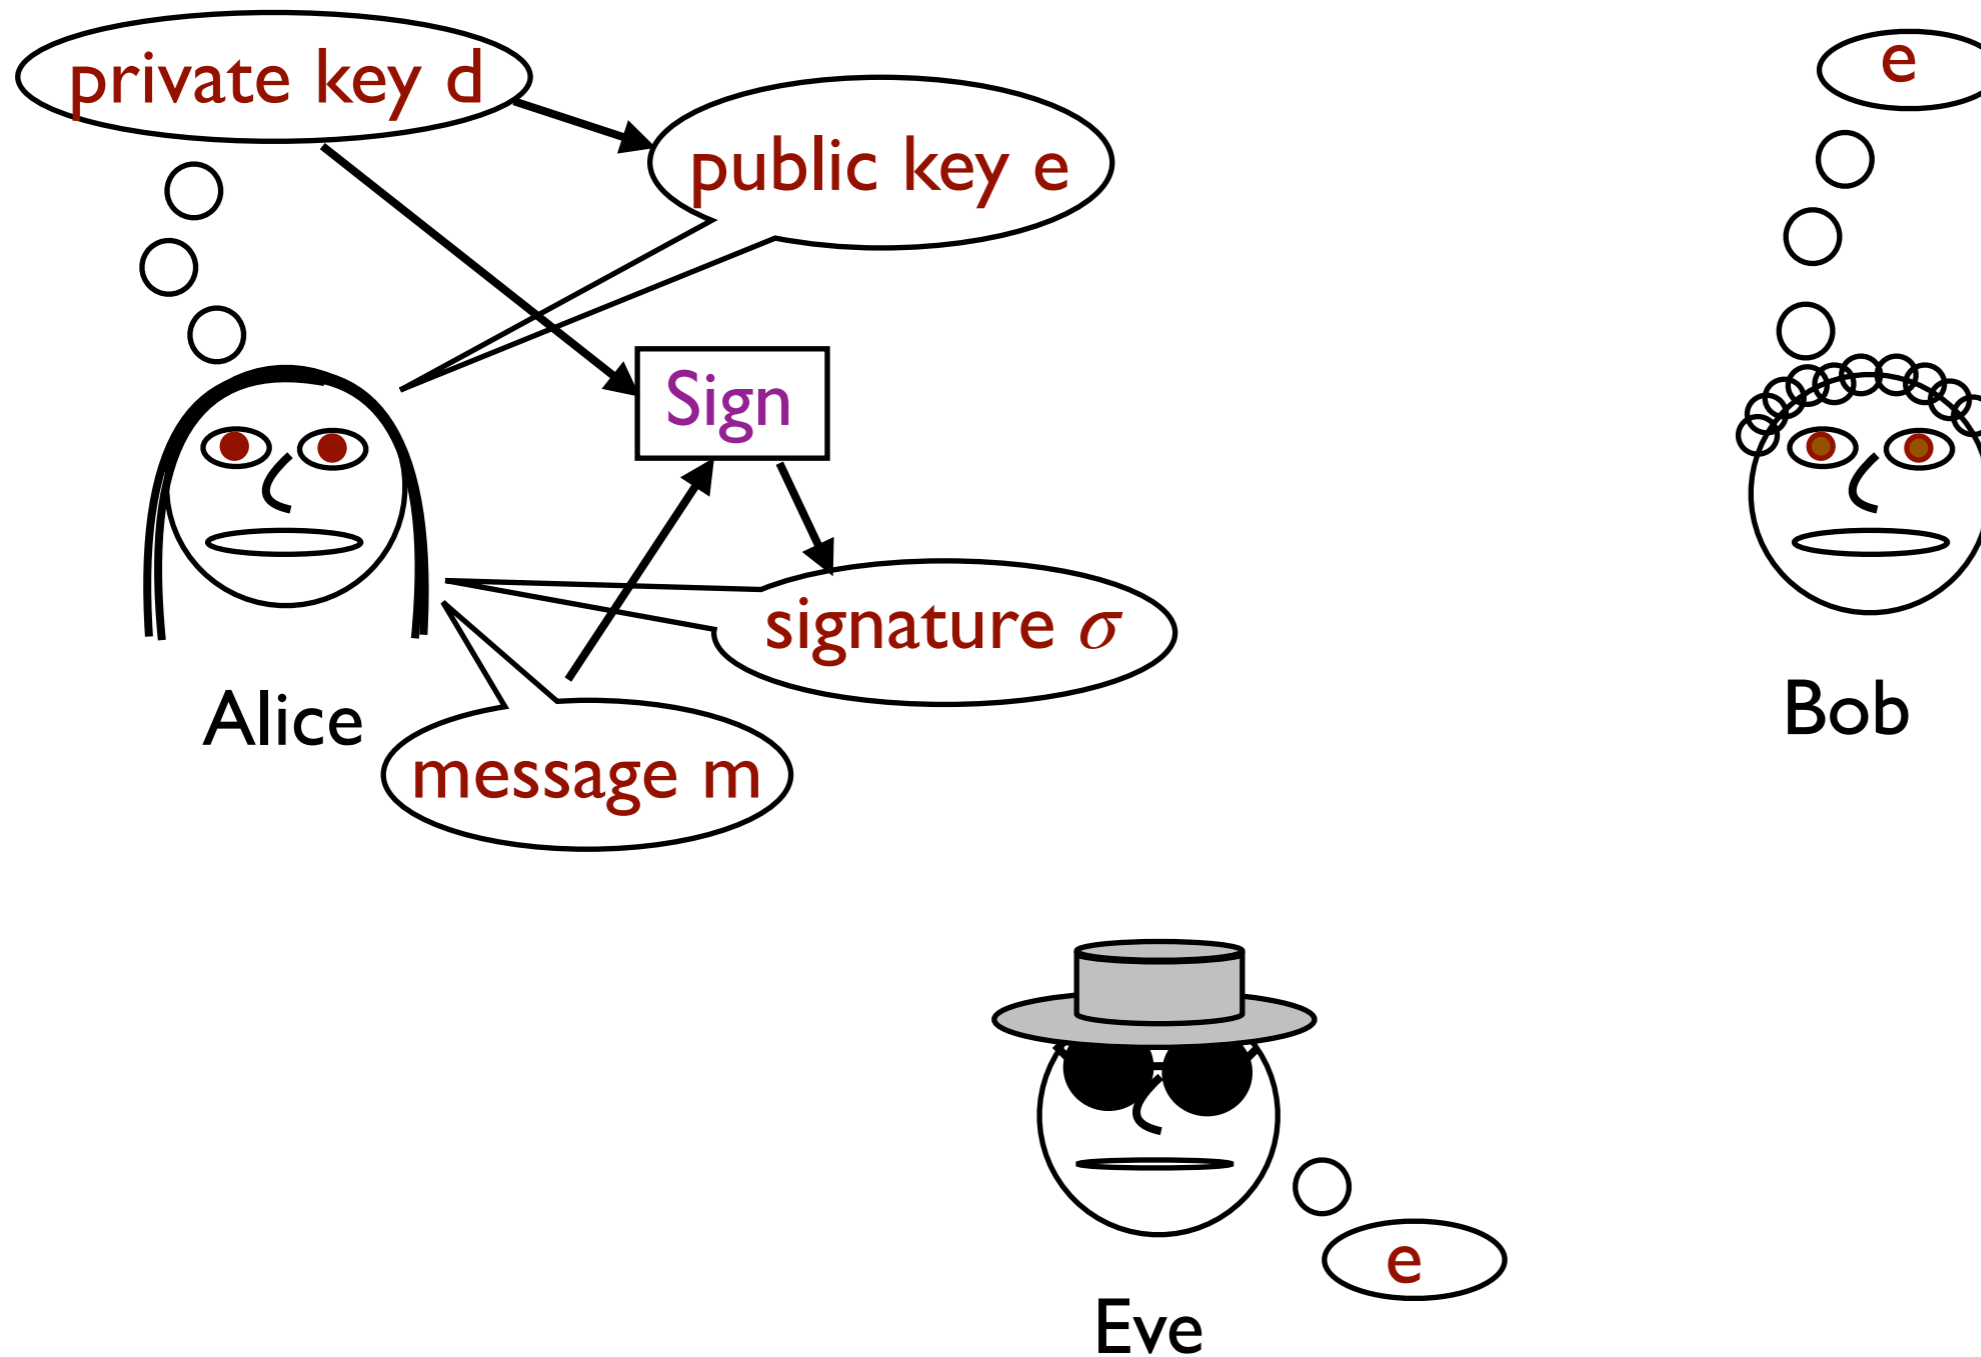
Digital Signatures

Digital signatures are a public key version of MACs.



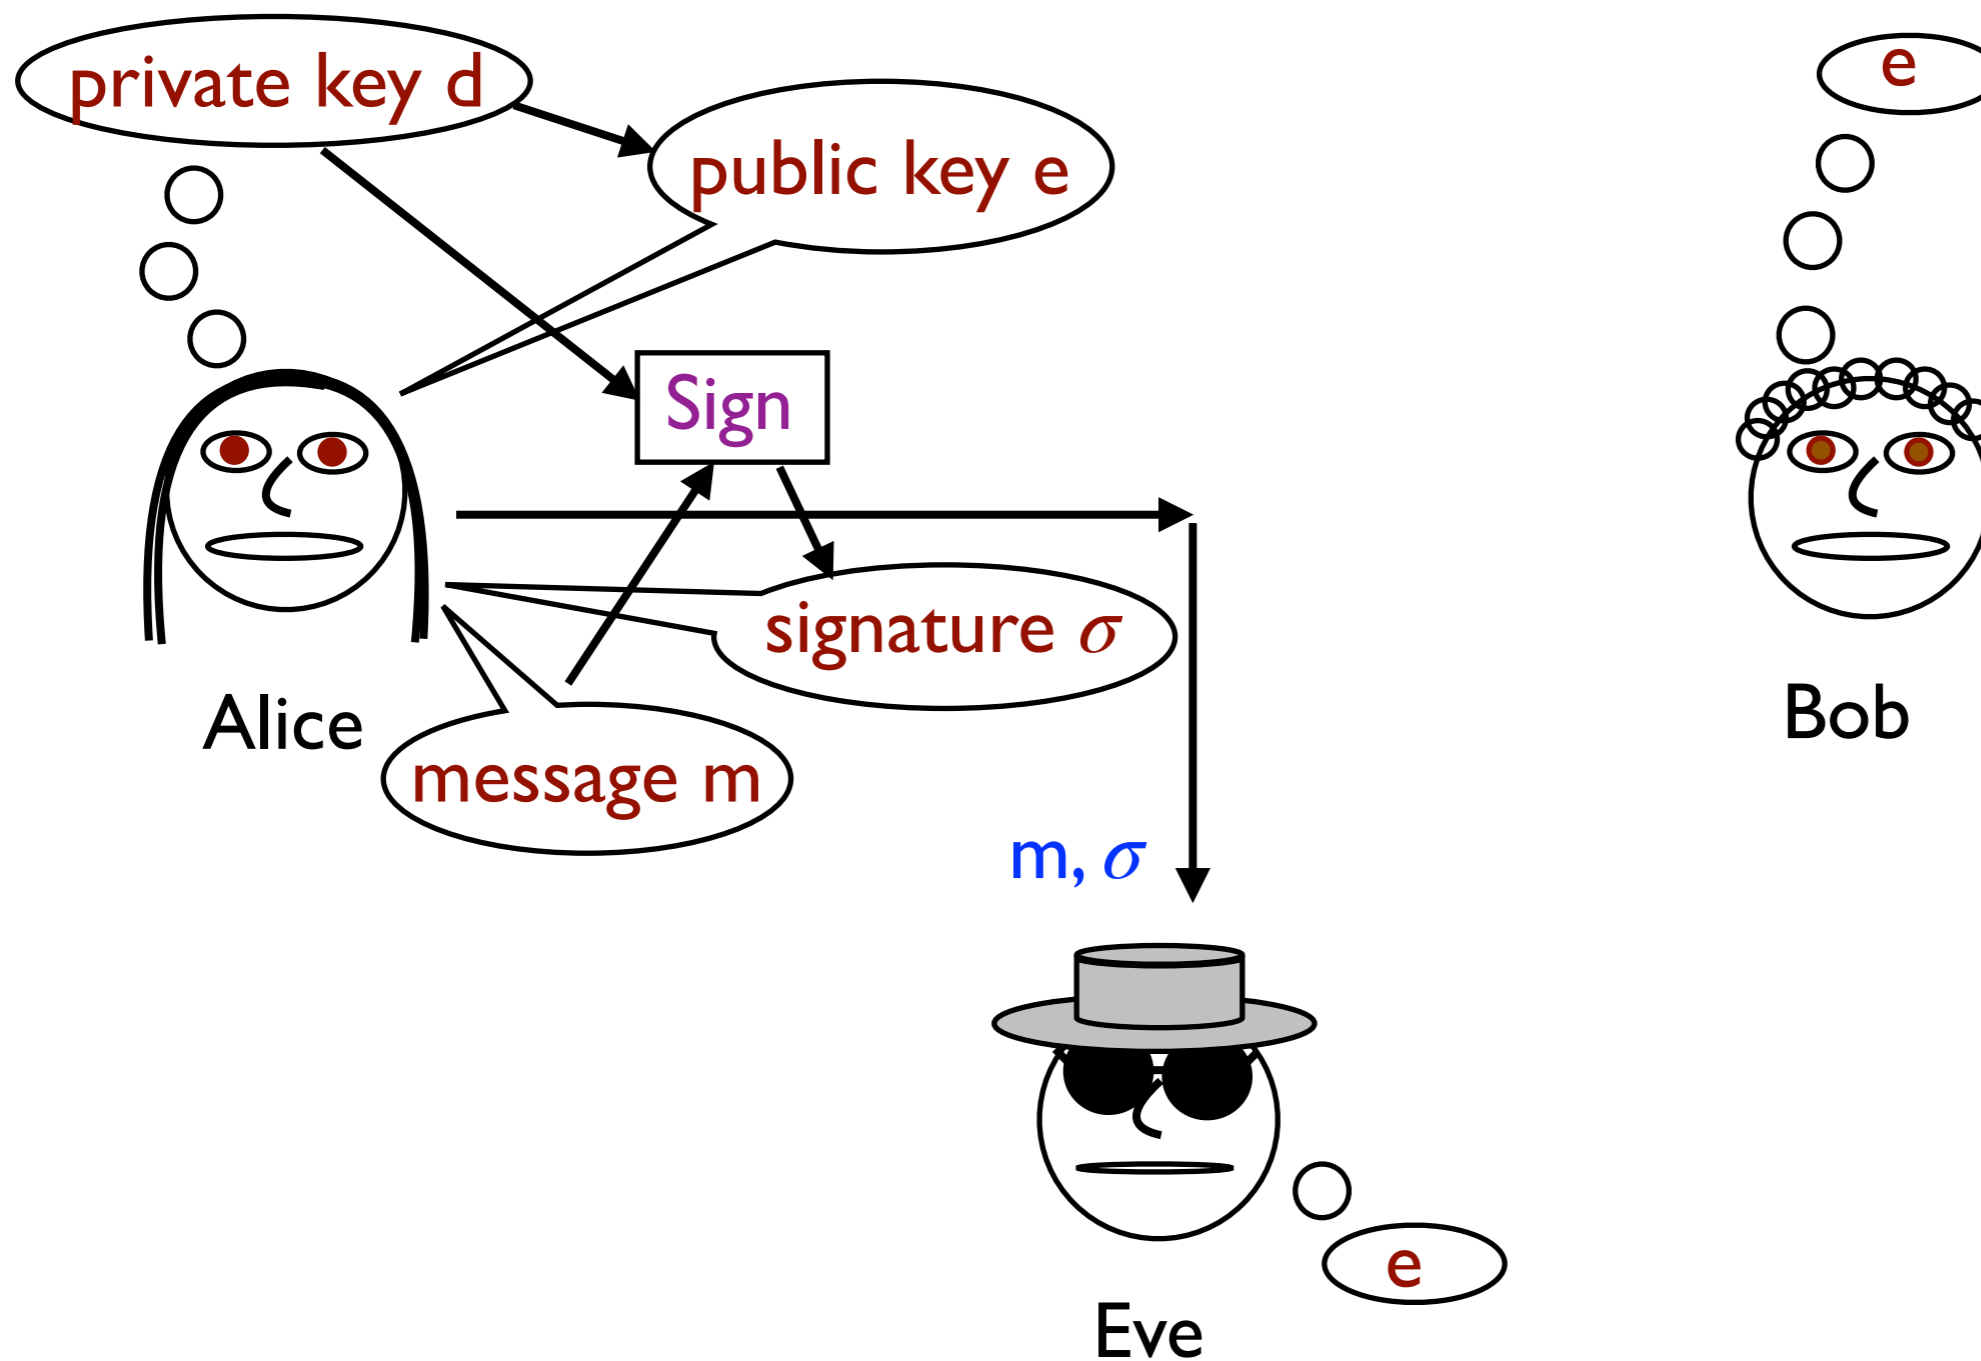
Digital Signatures

Digital signatures are a public key version of MACs.



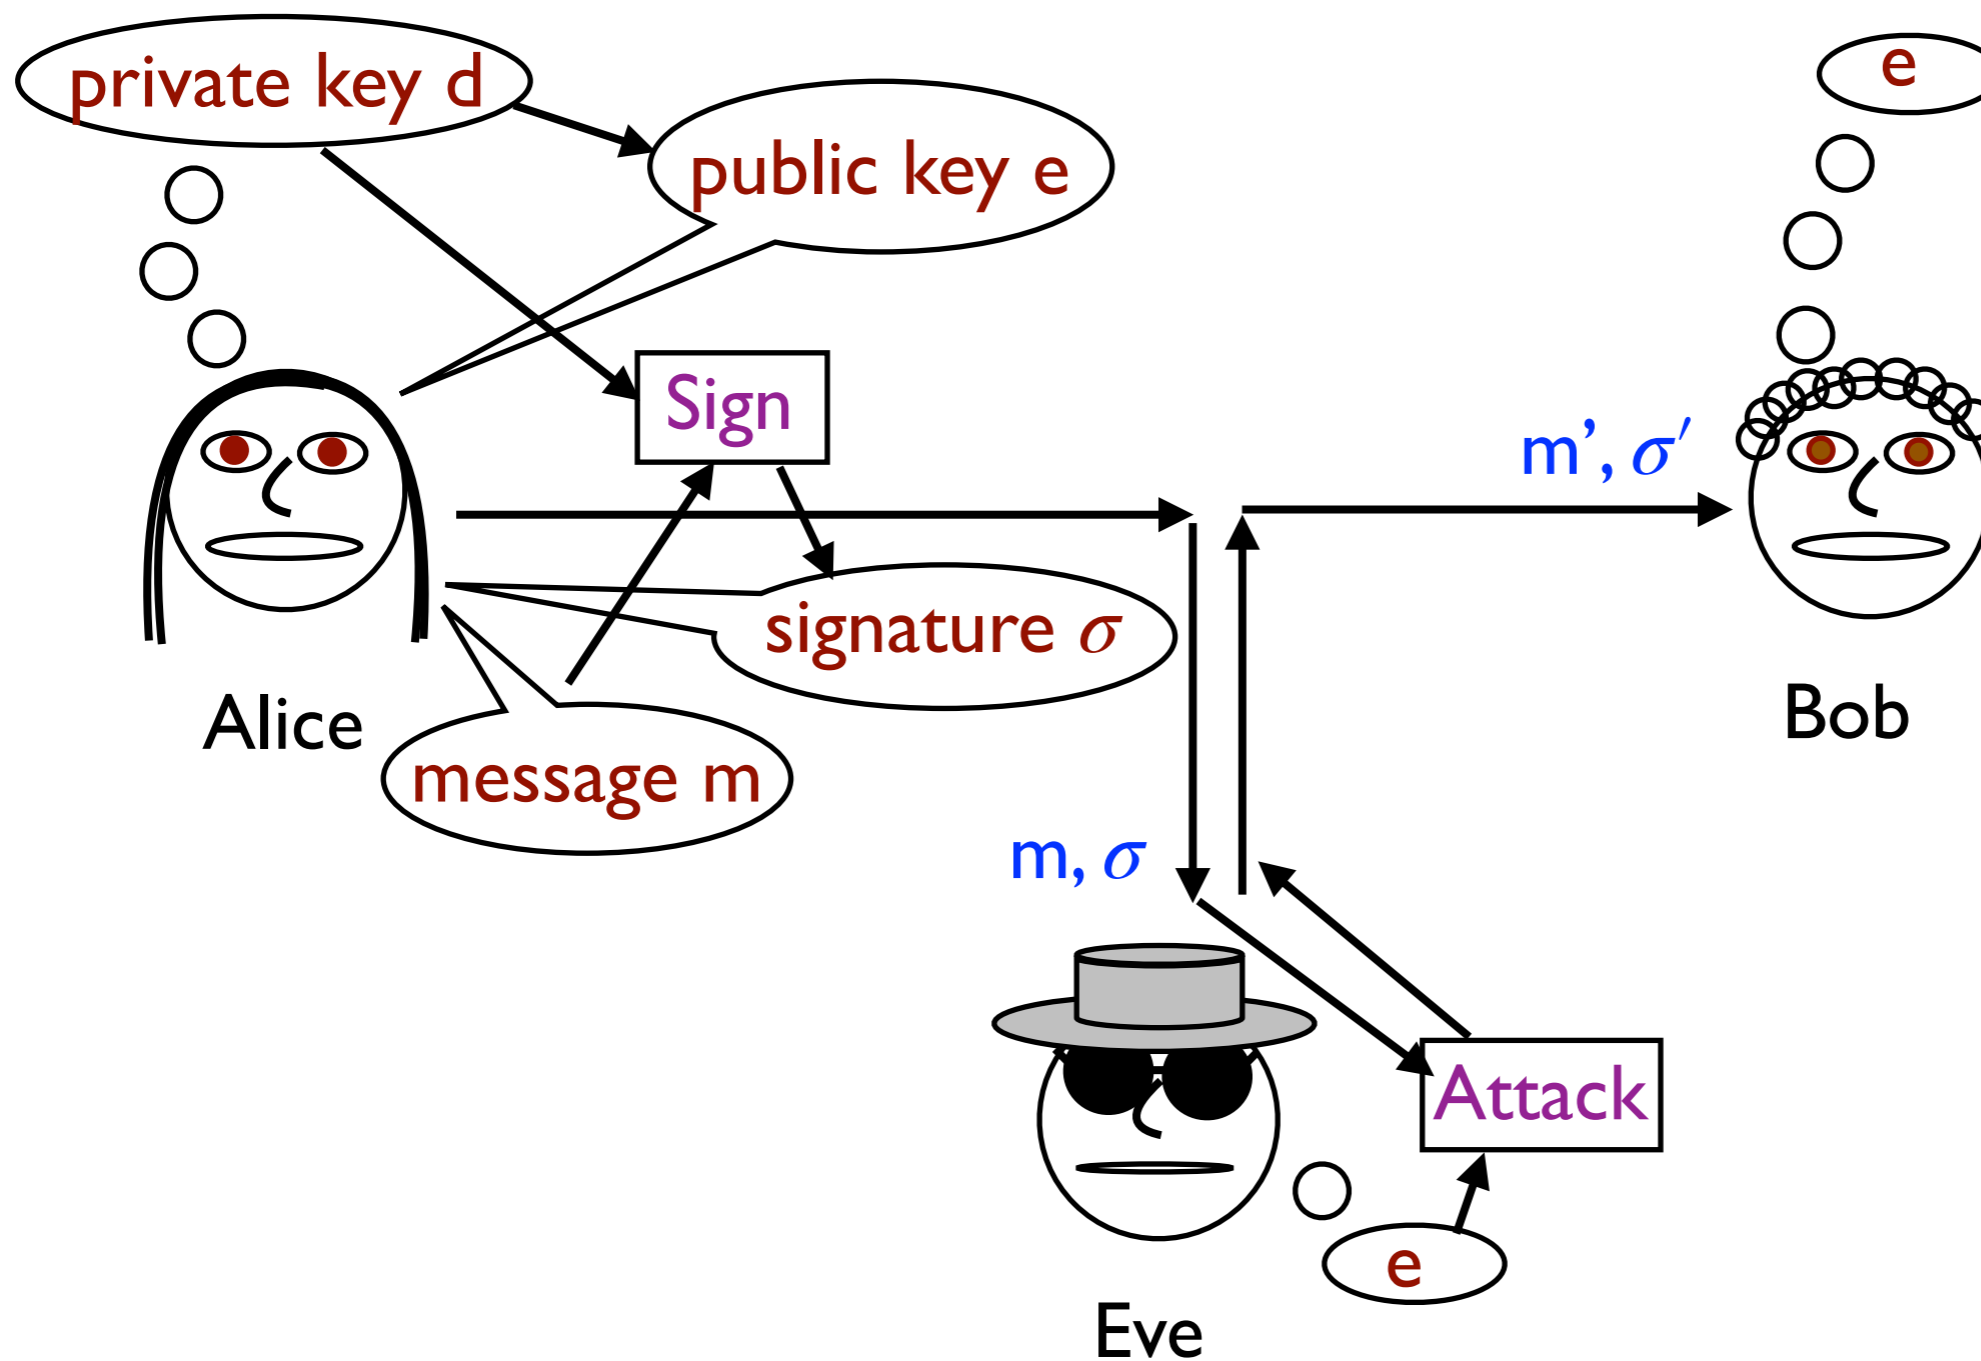
Digital Signatures

Digital signatures are a public key version of MACs.



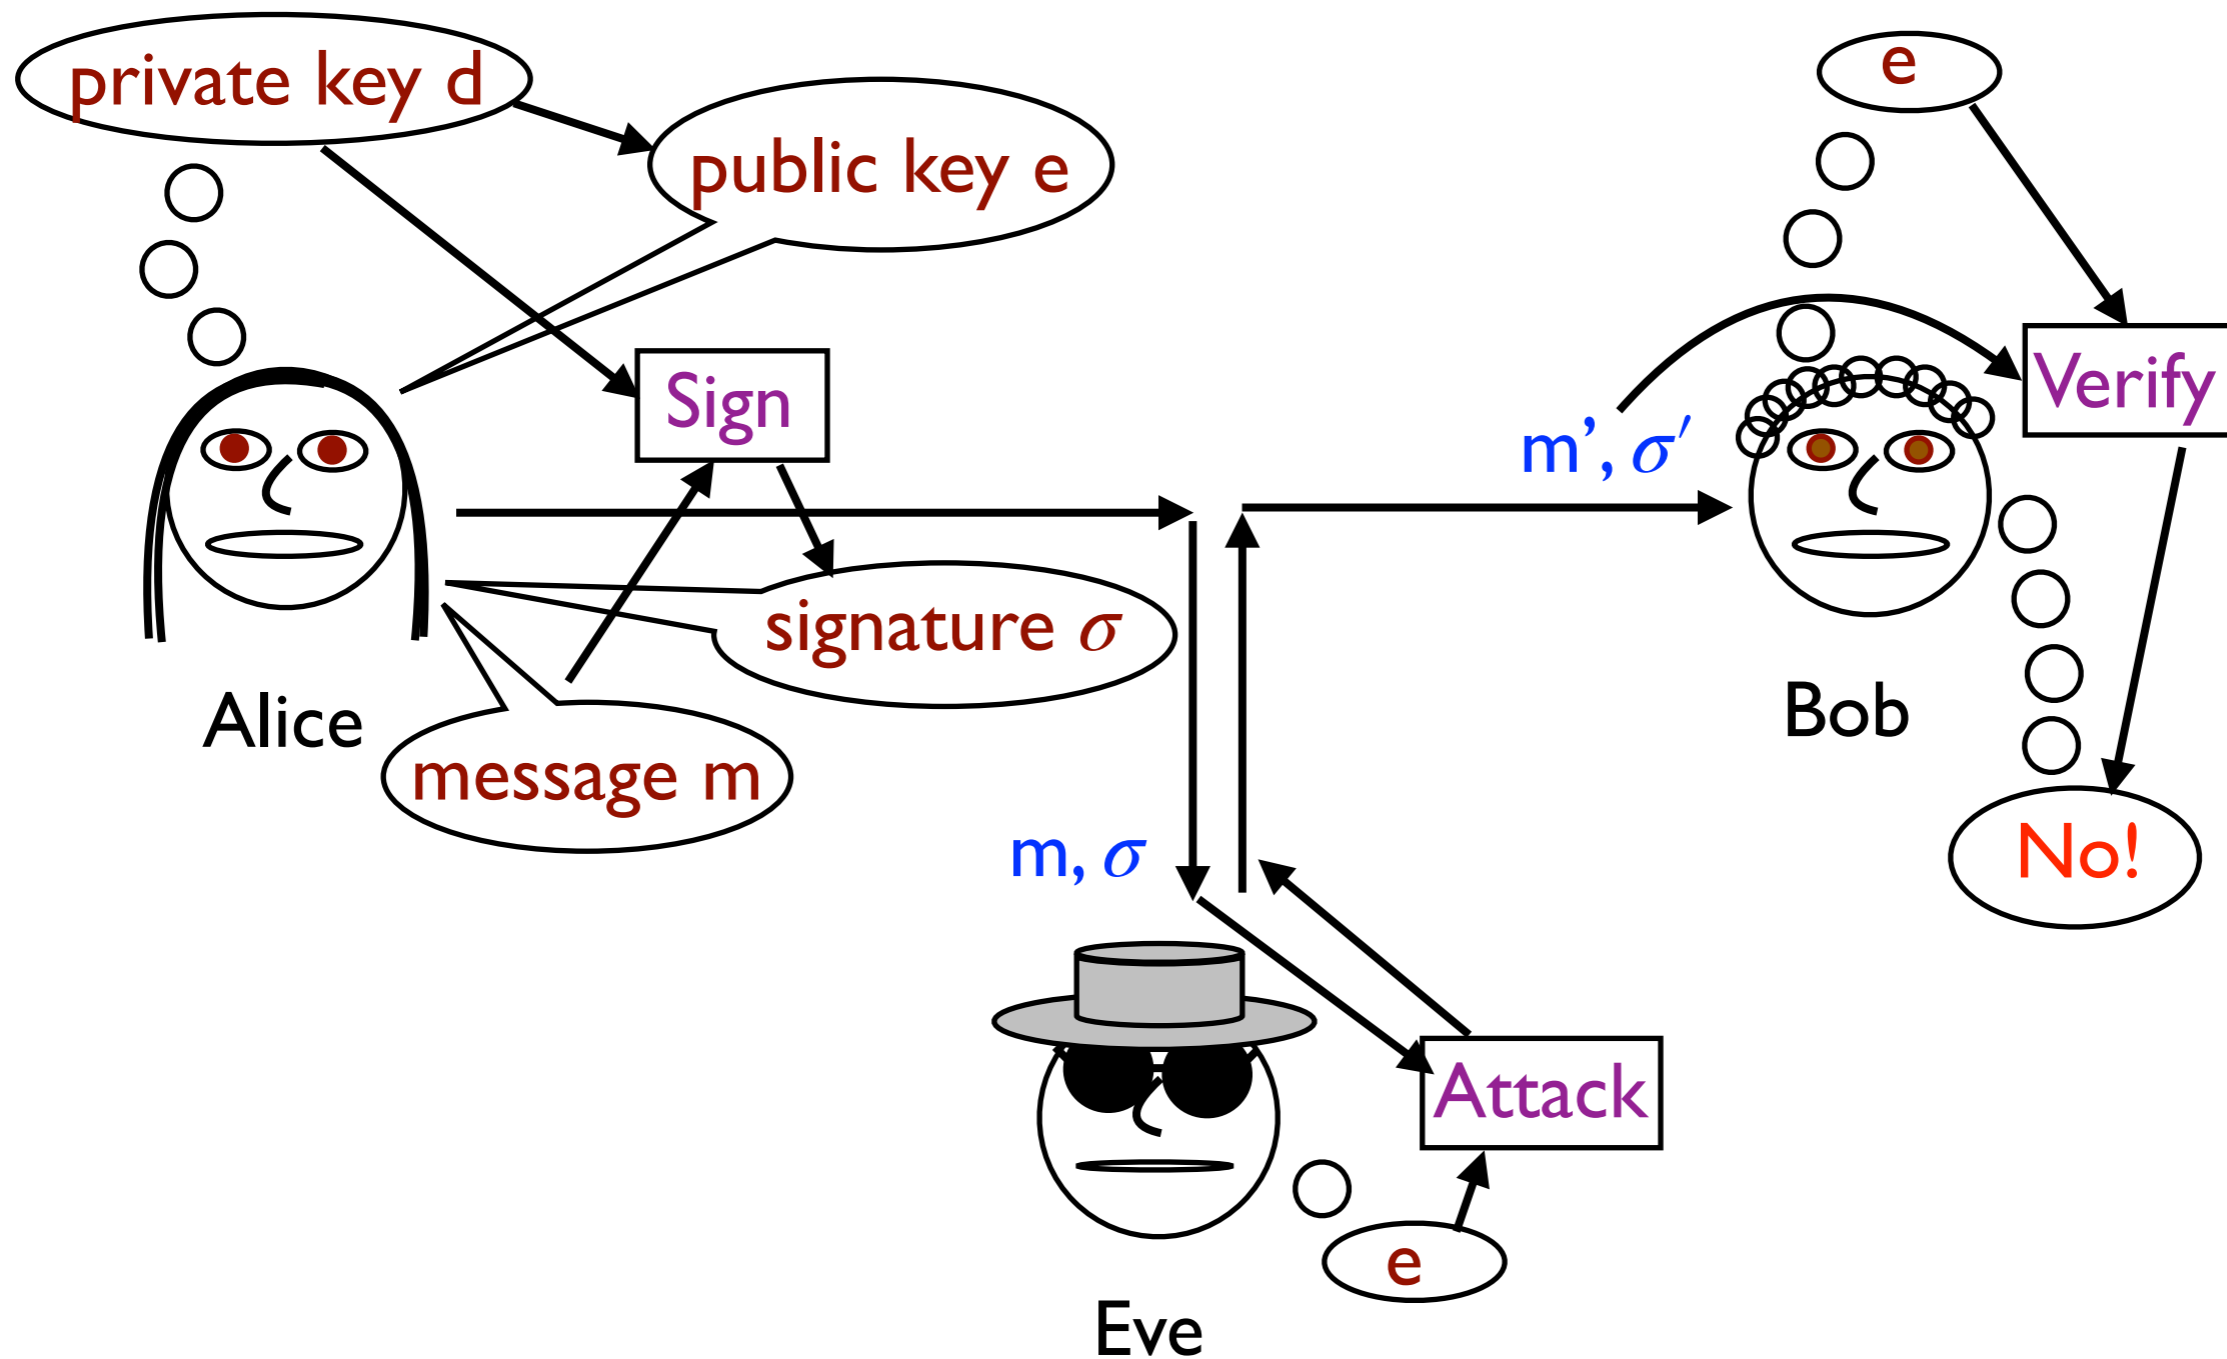
Digital Signatures

Digital signatures are a public key version of MACs.



Digital Signatures

Digital signatures are a public key version of MACs.



Digital Signature Definition

Definition: A **digital signature** is a set of three probabilistic polynomial-time algorithms (**Gen**, **Sign**, **Vrfy**):

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

Sign is the **signing algorithm**. It takes as input the private key **d** and a **message** $m \in \{0,1\}^*$ and outputs a **signature** $\sigma \in \{0,1\}^*$.

Vrfy is the **verification algorithm**. It takes as input the public key **e** and (m, σ) and outputs “**valid**” or “**invalid**.”

The digital signature scheme is **correct** if

$$\text{Vrfy}(e, m, \text{Sign}(d, m)) = \text{valid}$$

Digital Signature Definition

Definition: A **digital signature** is a set of three probabilistic polynomial-time algorithms (**Gen**, **Sign**, **Vrfy**):

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

Sign is the **signing algorithm**. It takes as input the private key **d** and a **message** $m \in \{0,1\}^*$ and outputs a **signature** $\sigma \in \{0,1\}^*$.


Vrfy is the **verification algorithm**. It takes as input the public key **e** and (m, σ) and outputs “**valid**” or “**invalid**.”

The digital signature scheme is **correct** if

$$\text{Vrfy}(e, m, \text{Sign}(d, m)) = \text{valid}$$


Note: Unlike a MAC, **Vrfy** cannot just generate a new signature, since that would require the private key.

Transferability



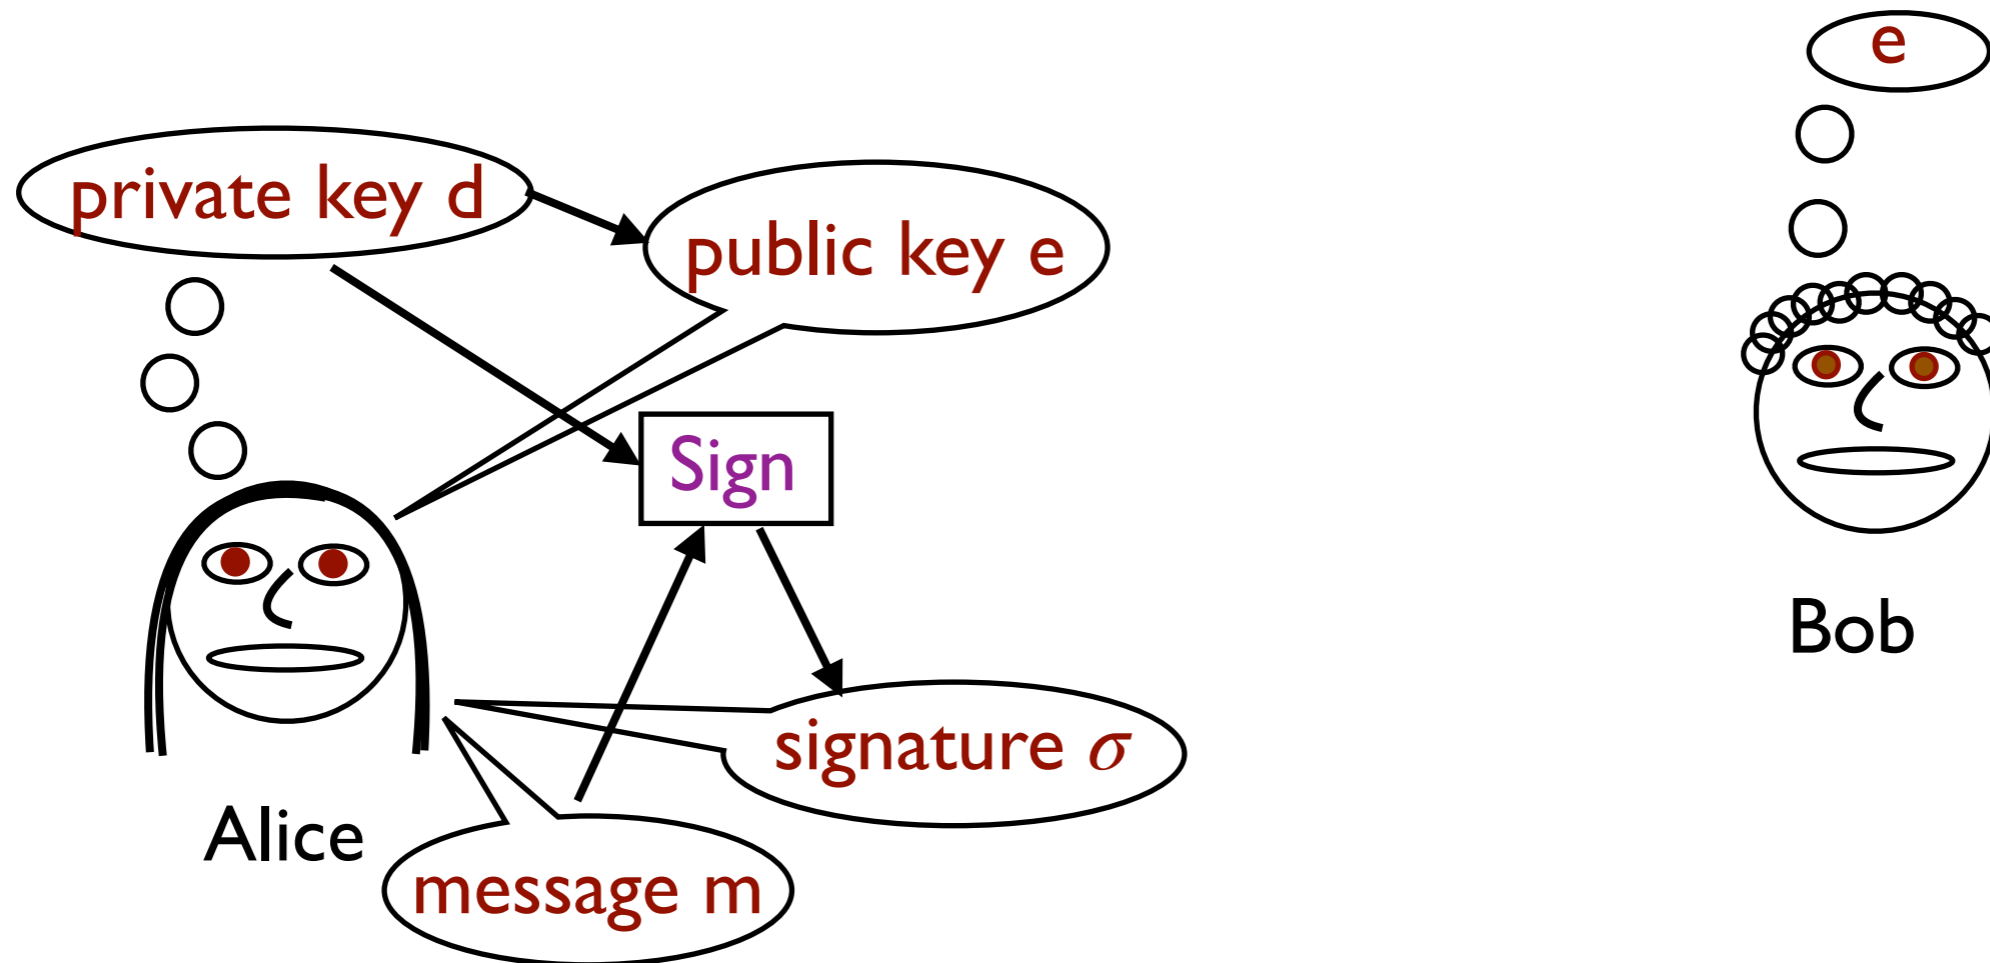
Because signatures use a public key, they are **transferable** between recipients.

Transferability



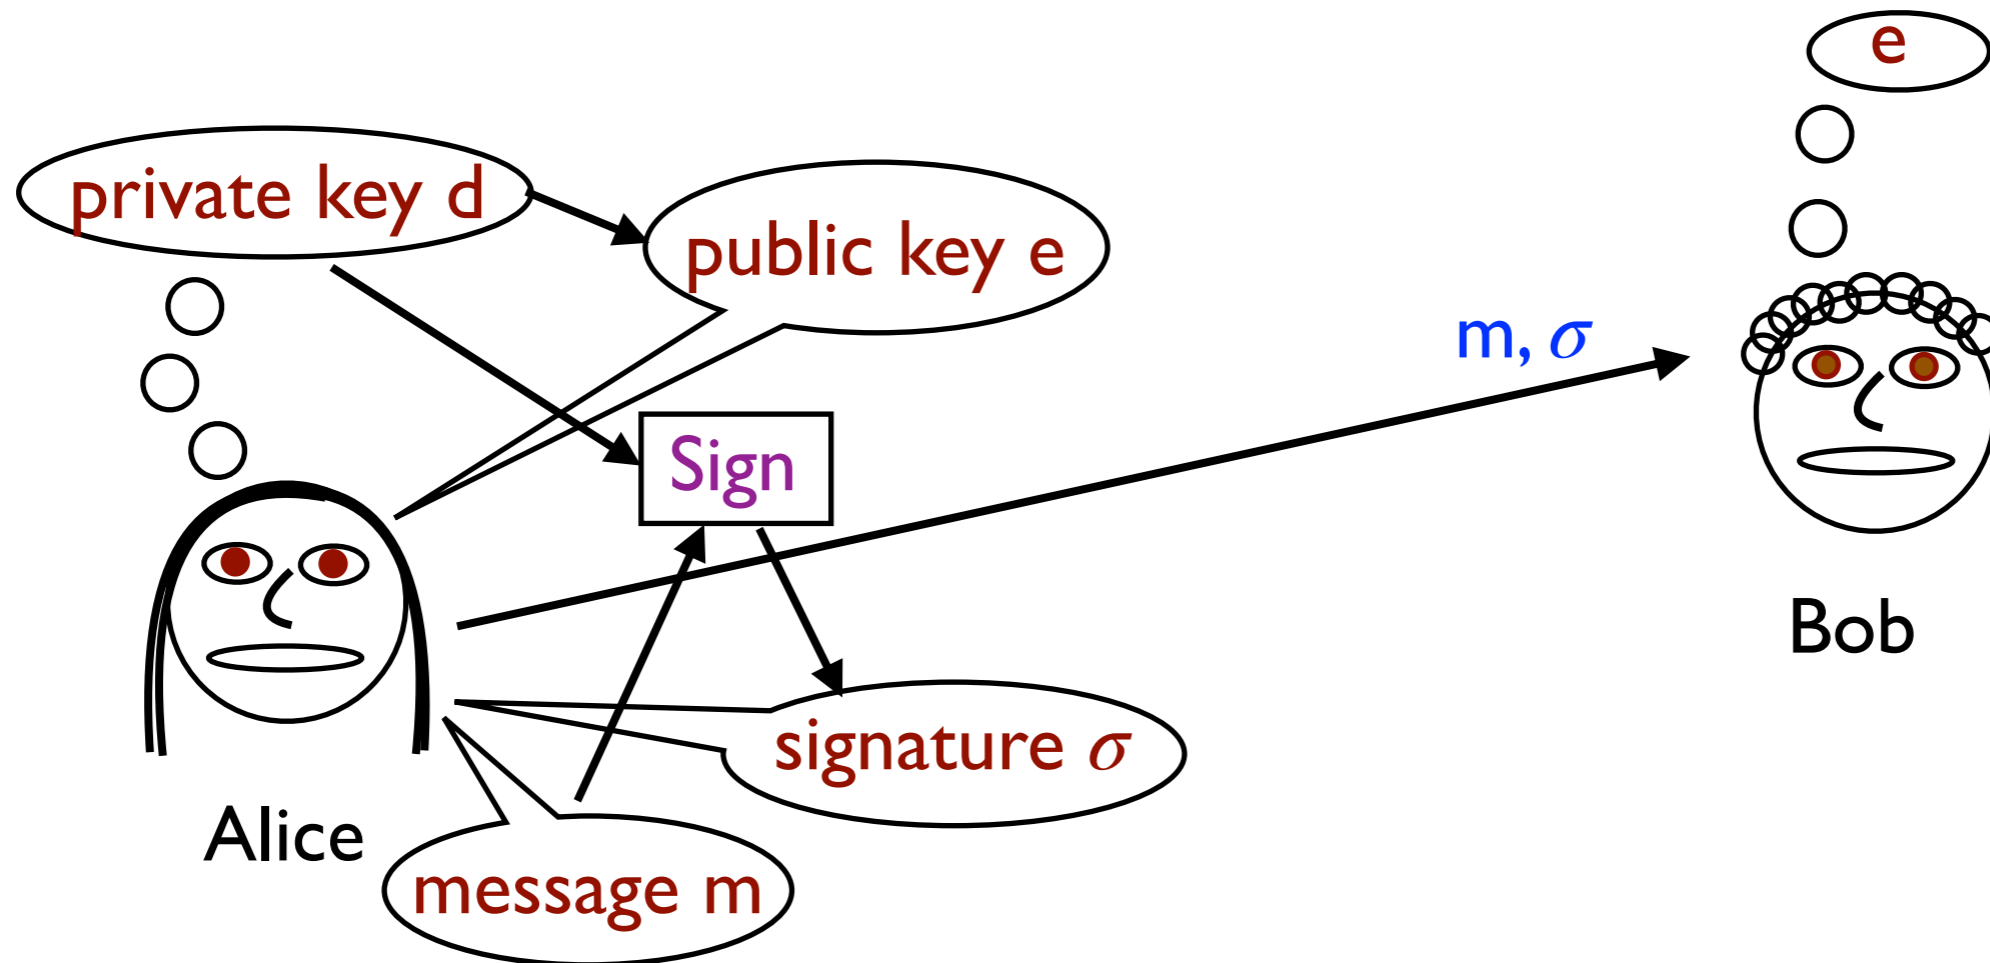
Because signatures use a public key, they are **transferable** between recipients.

Transferability



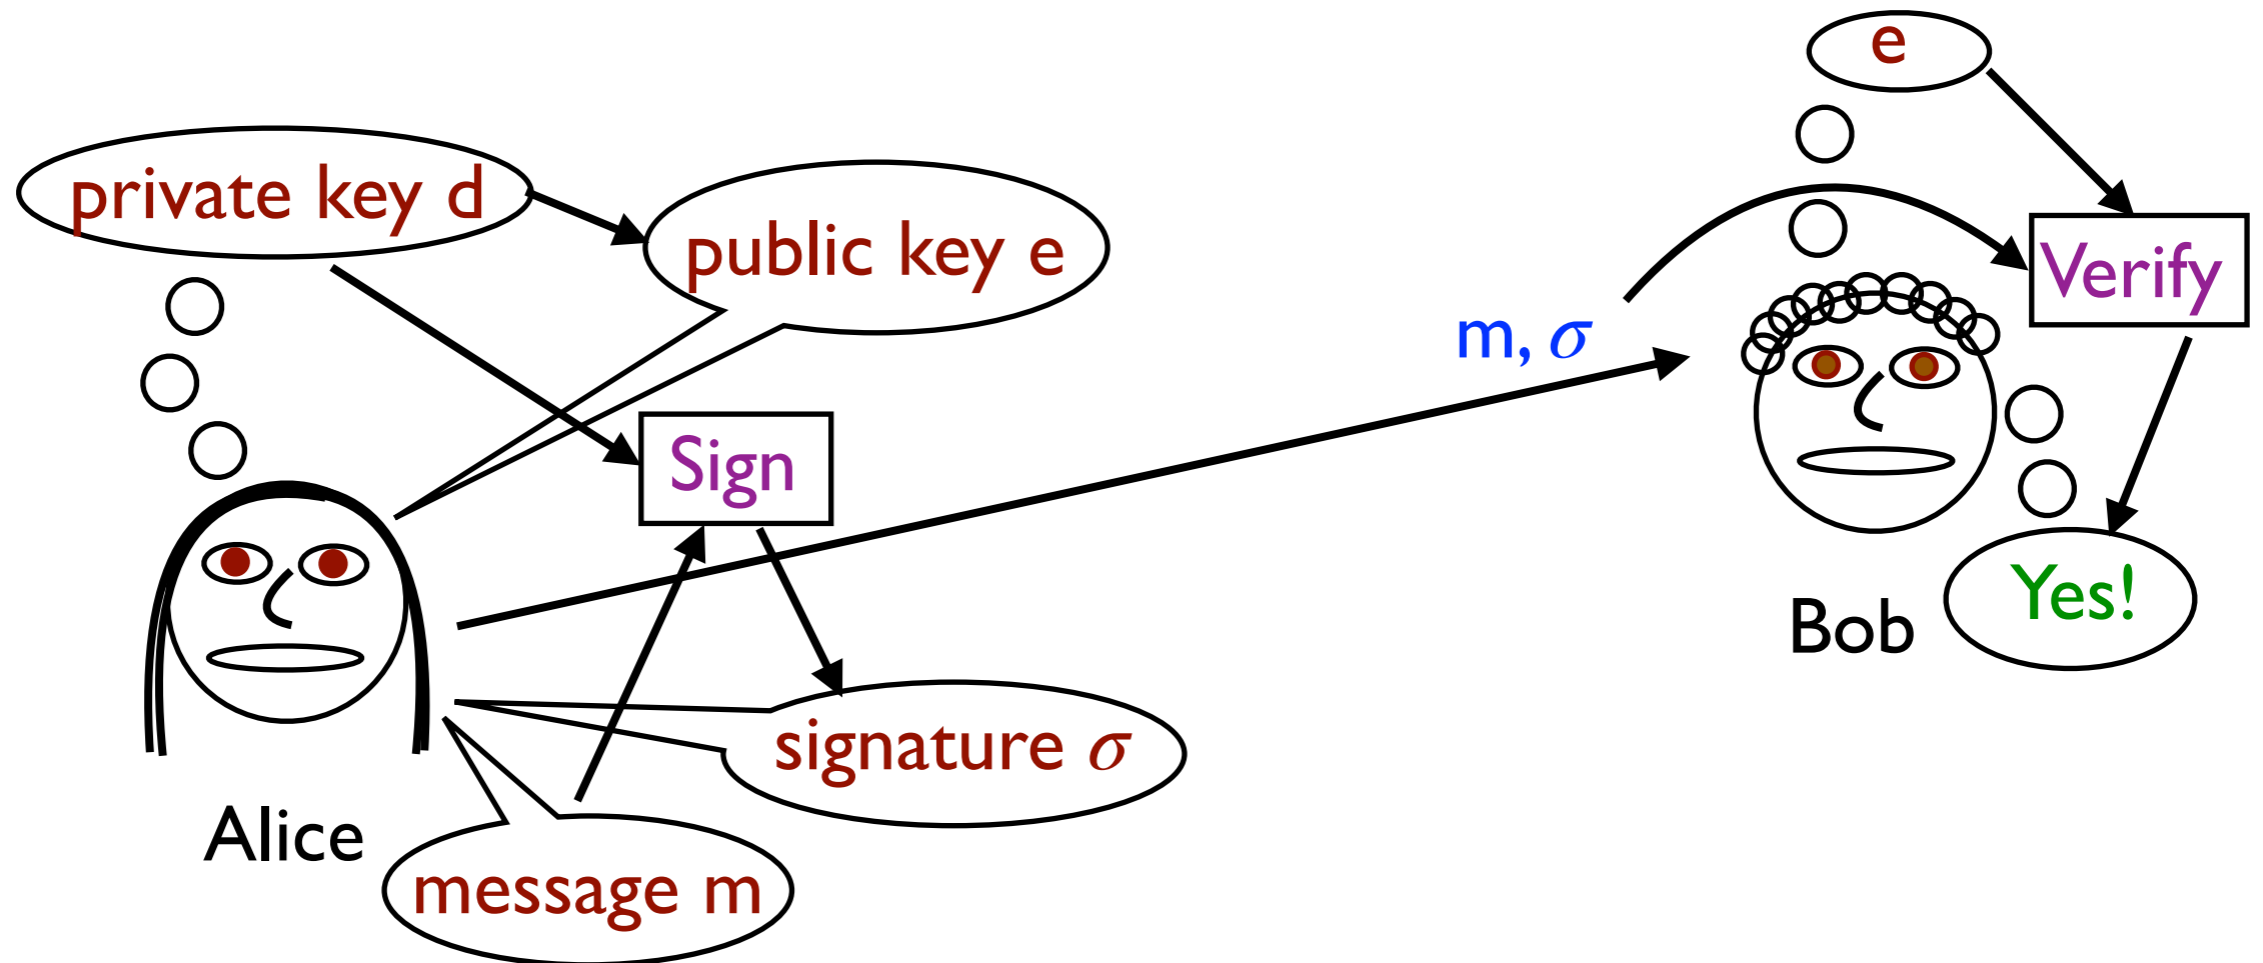
Because signatures use a public key, they are **transferable** between recipients.

Transferability



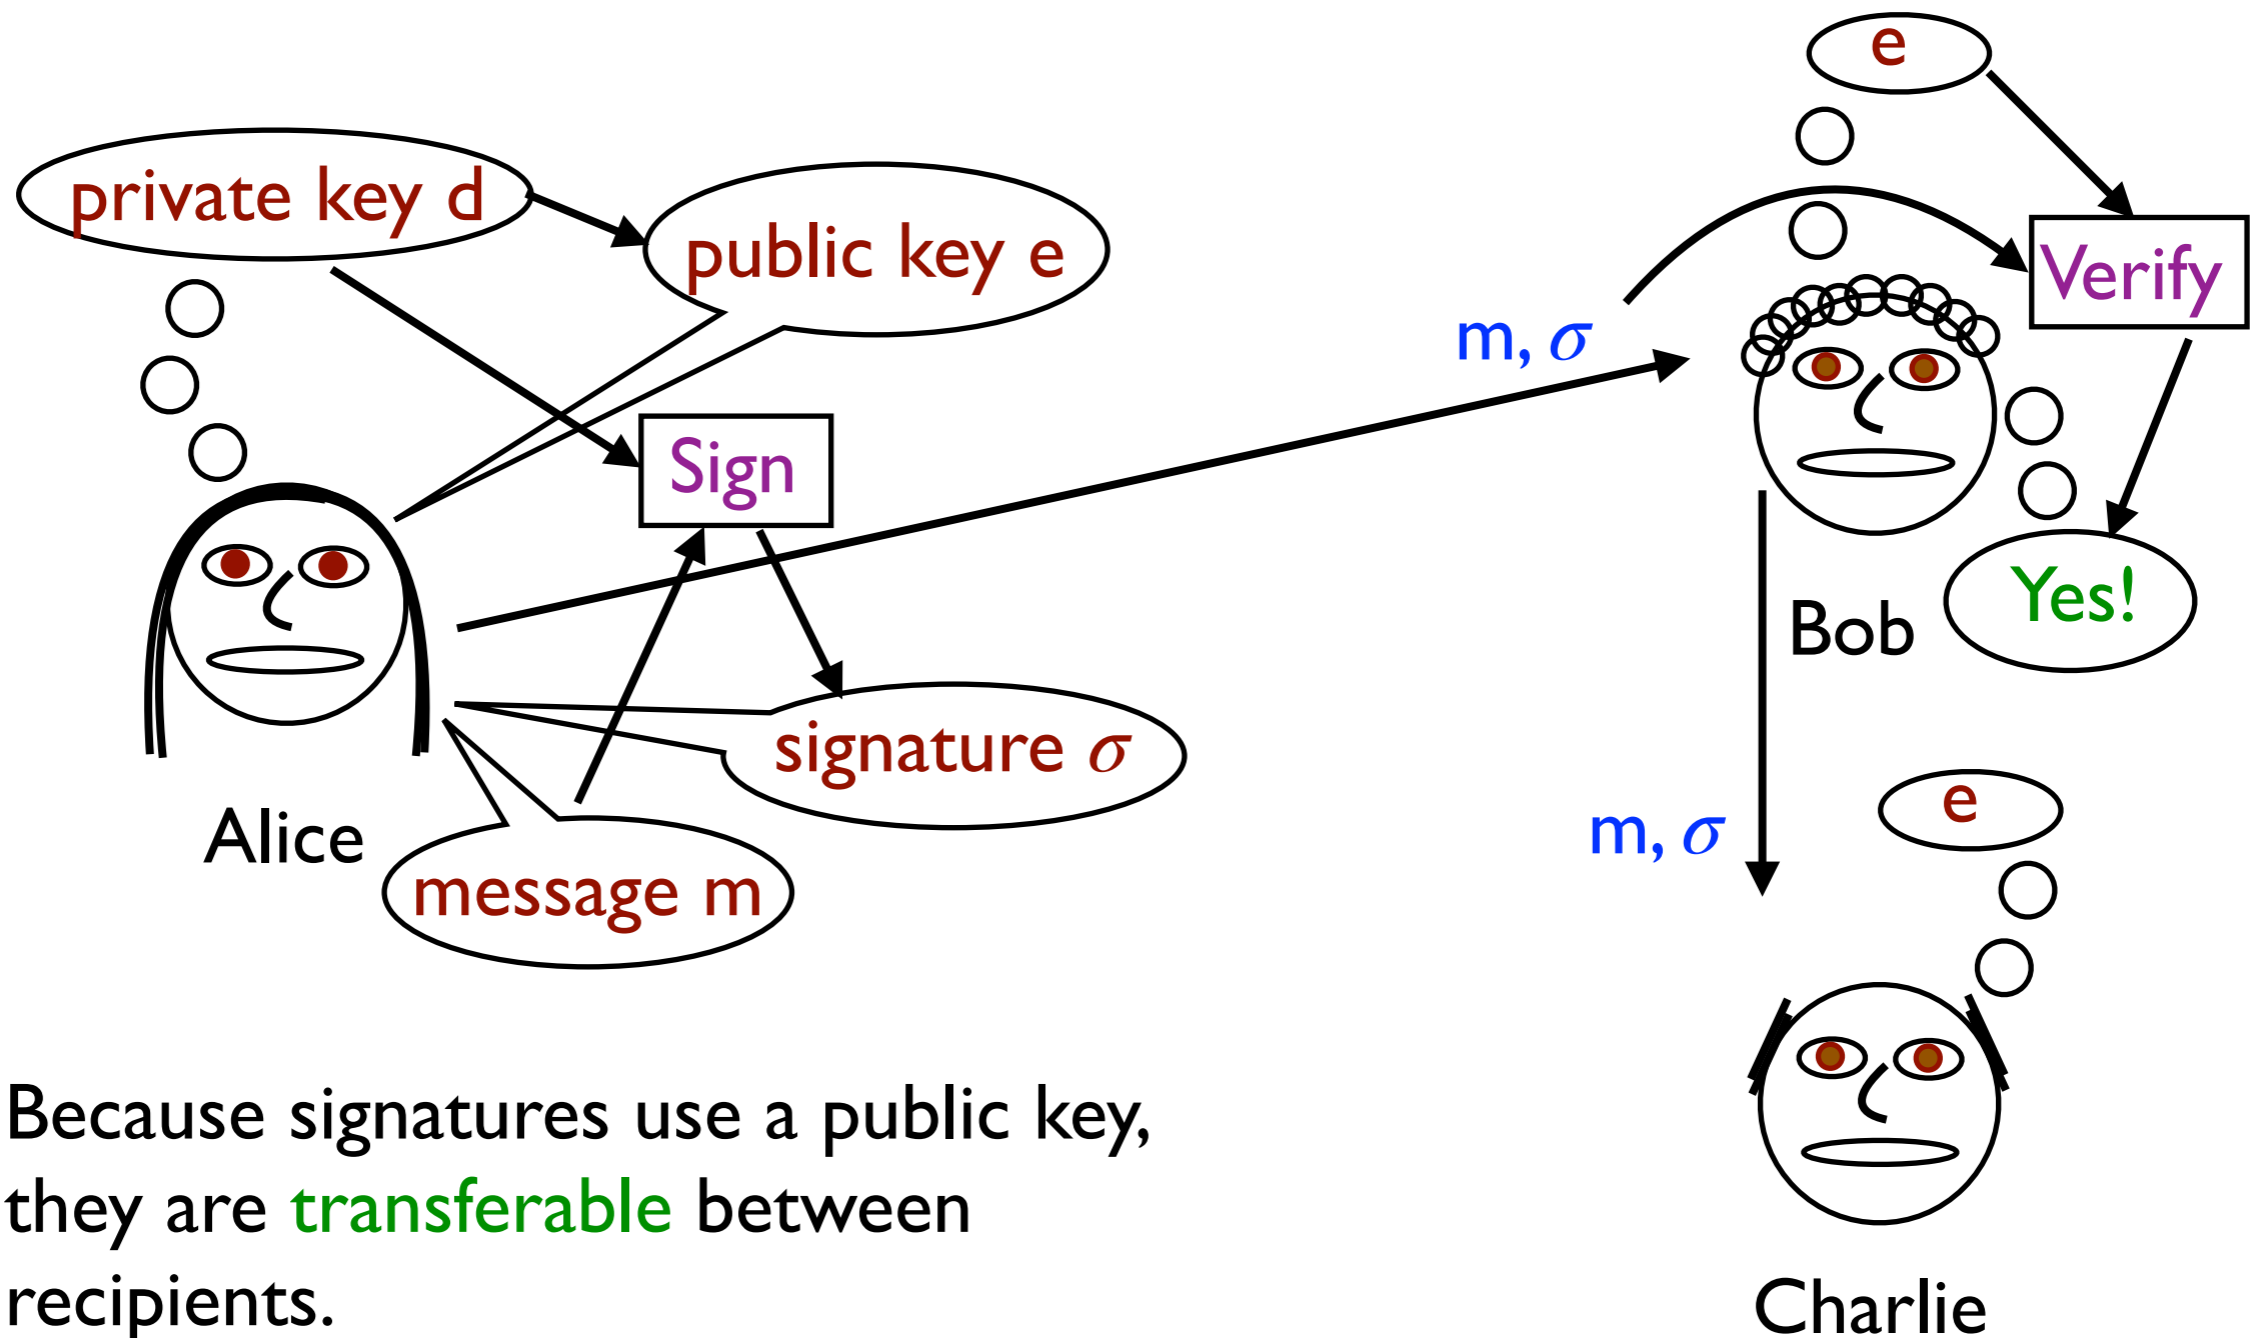
Because signatures use a public key, they are **transferable** between recipients.

Transferability



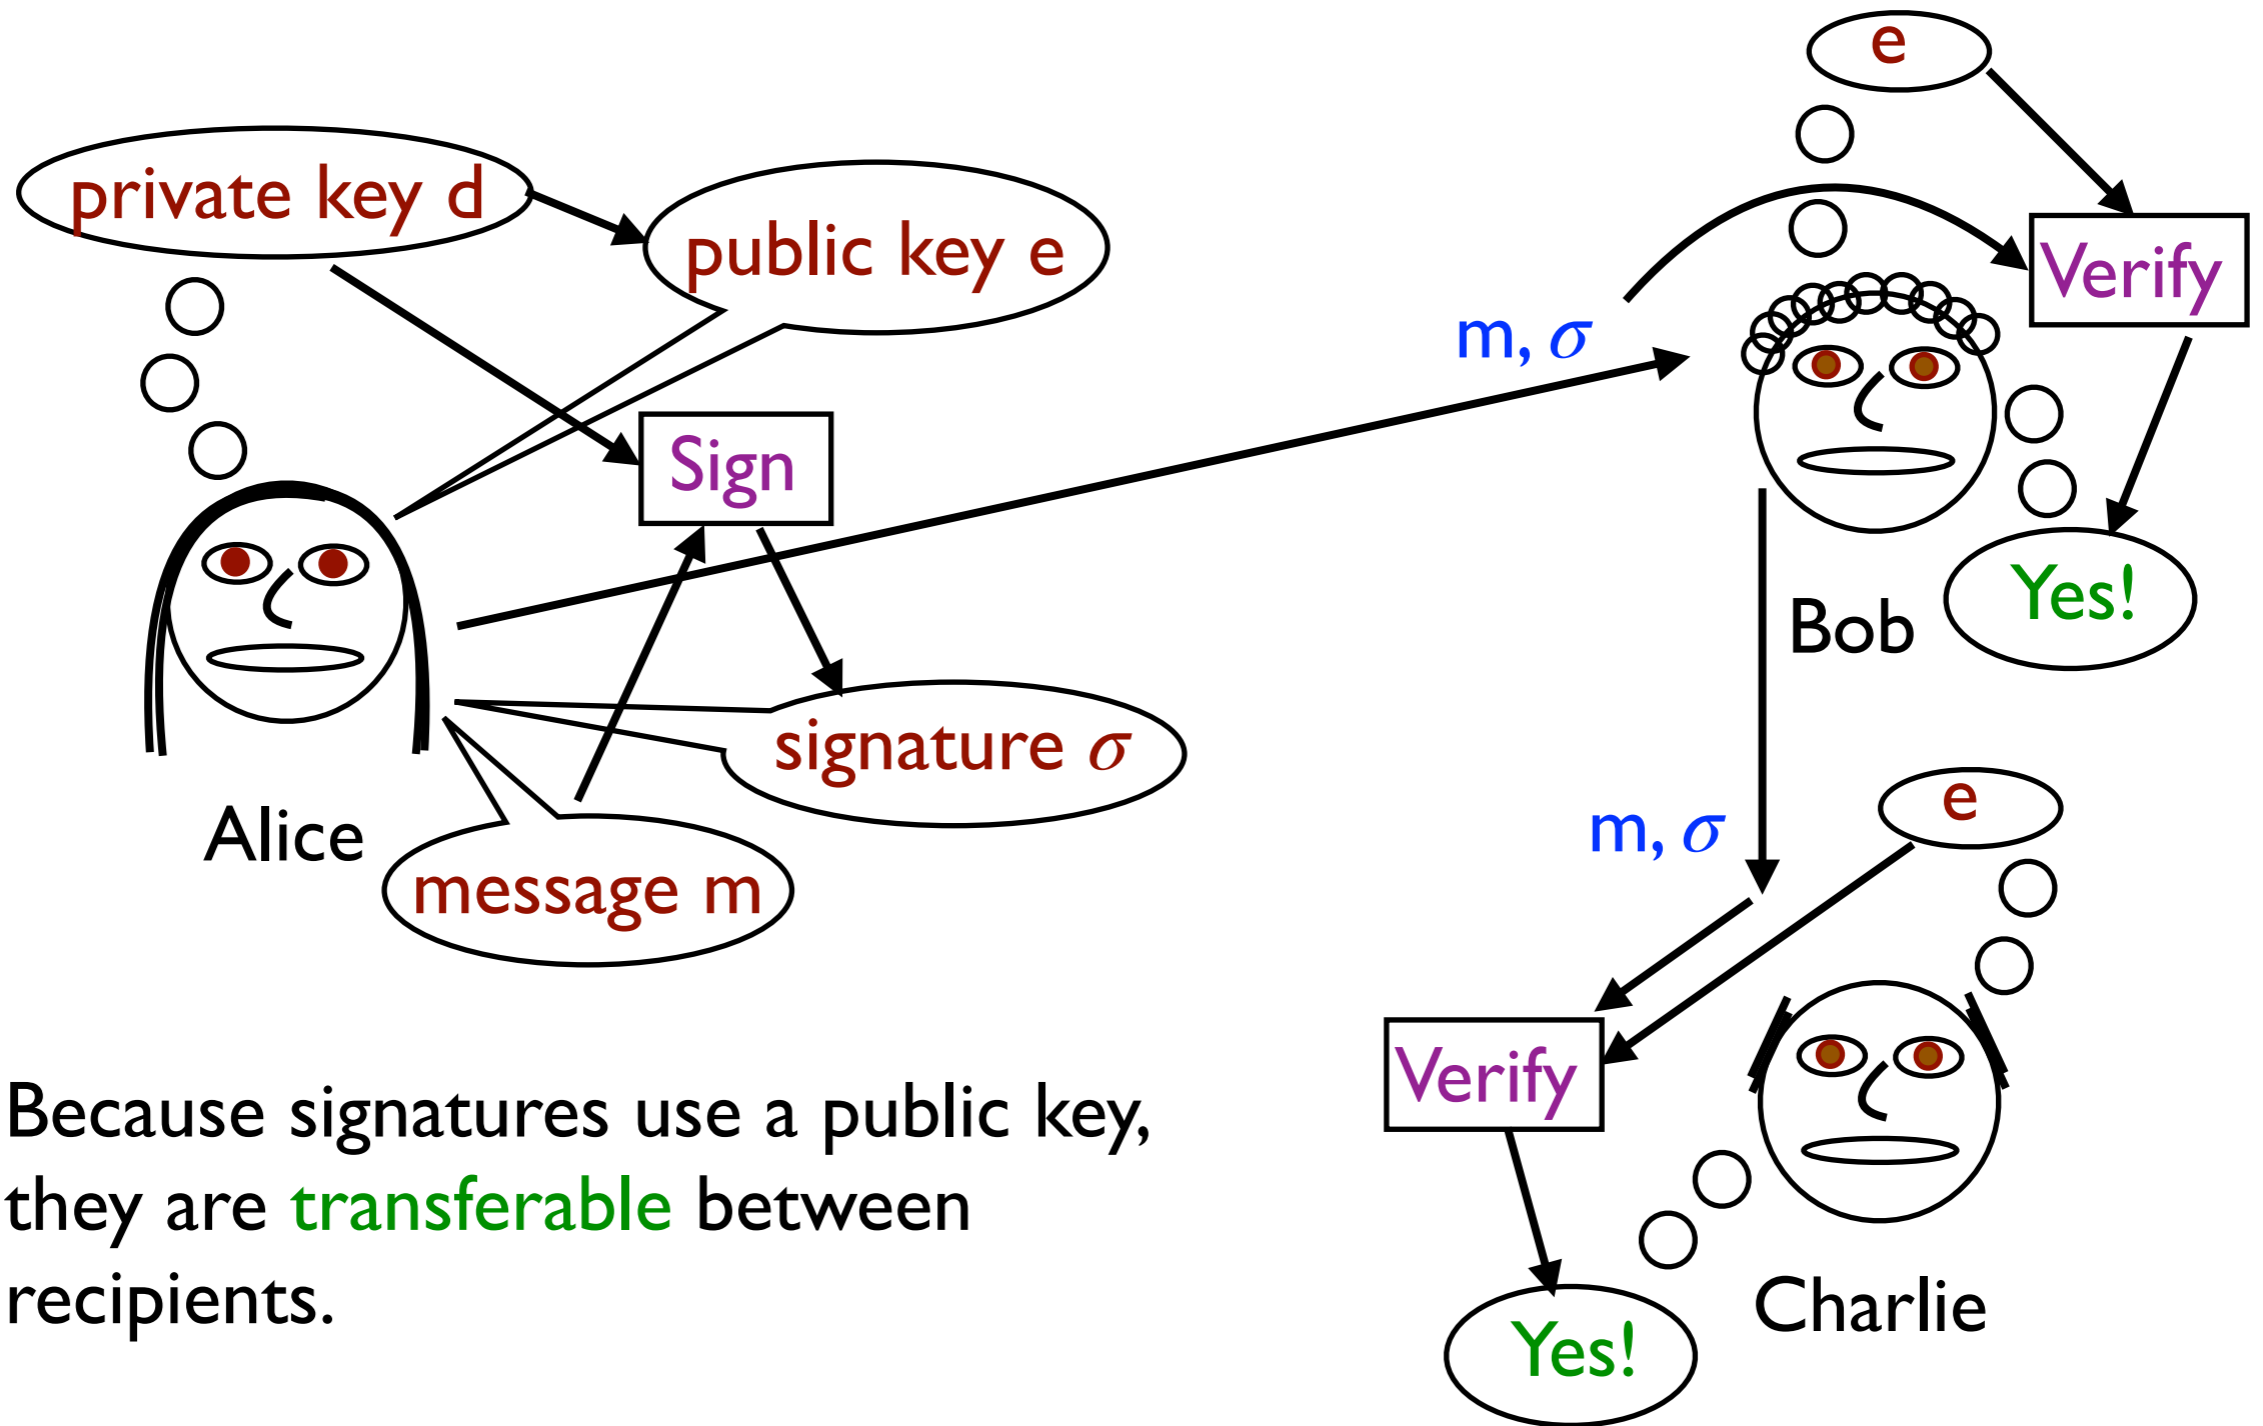
Because signatures use a public key, they are **transferable** between recipients.

Transferability



Because signatures use a public key, they are **transferable** between recipients.

Transferability



Because signatures use a public key, they are **transferable** between recipients.

Digital Signature vs. MAC

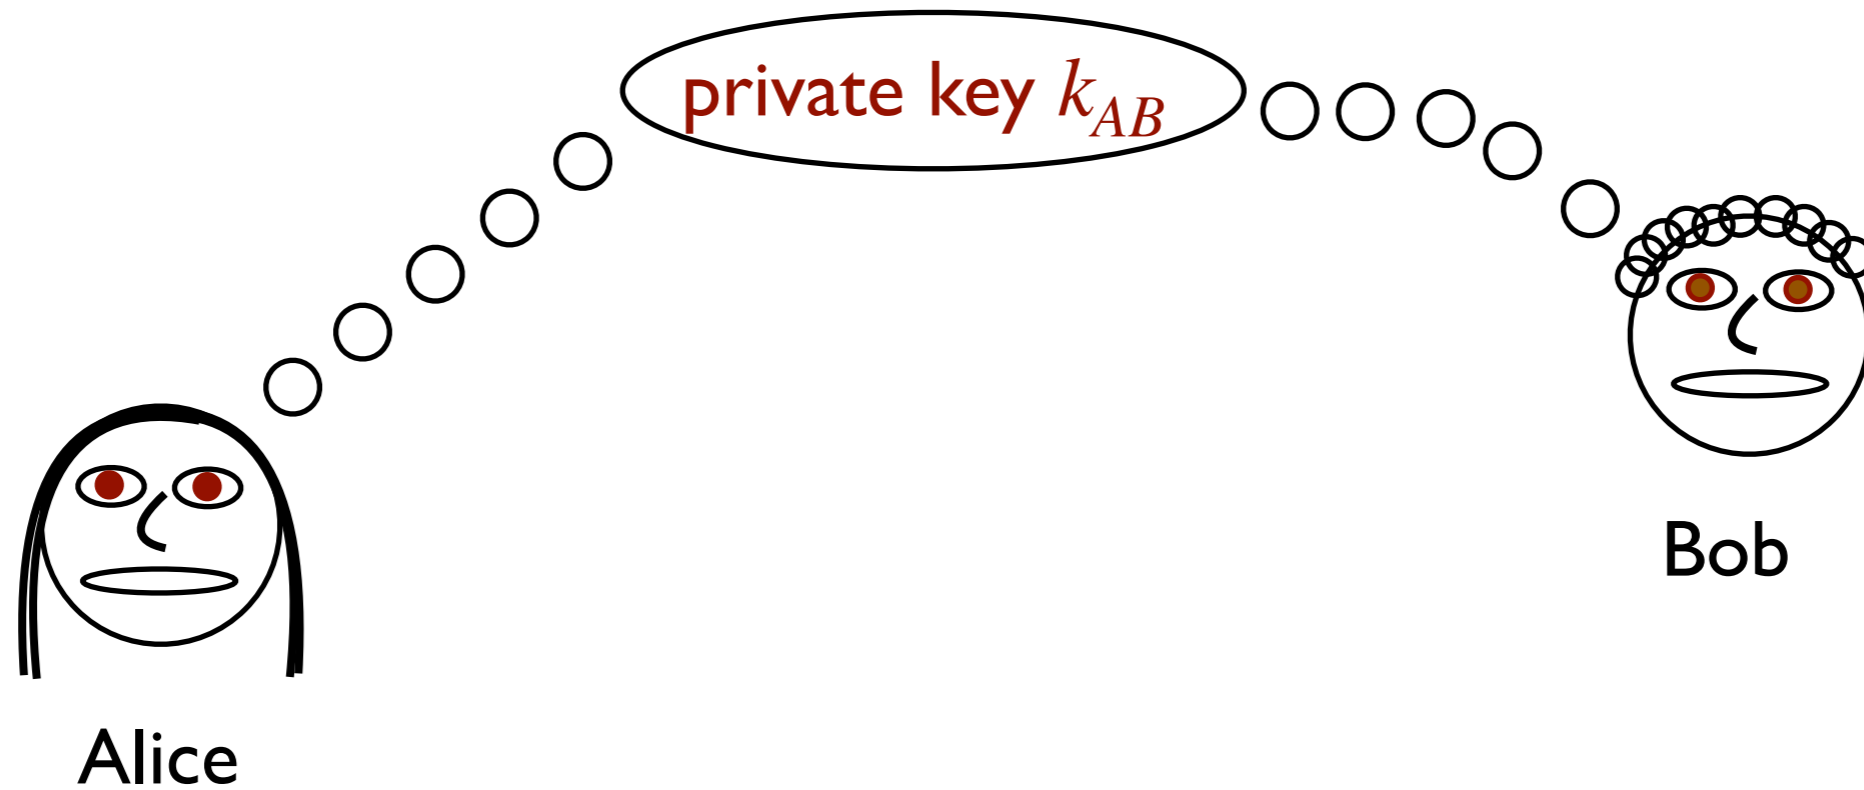
In a **MAC**, each prospective recipient must use a **different key**. Otherwise, any valid recipient can also forge messages.

This, in turn, means that to send the same message to many recipients, you have to use a different tag for each recipient.


In a **digital signature**, the **same public key** is used for everyone. That means that the same signature can be used for any number of recipients. All of them verify it using the same information.

- Digital signatures are more efficient when sending to many recipients.
- Digital signatures allow you to sign a message to someone you don't know (provided they have your public key).
- Transferability allows you to pass on a digitally signed message to someone else.
 - This is critical for certificate authorities and public key infrastructure.

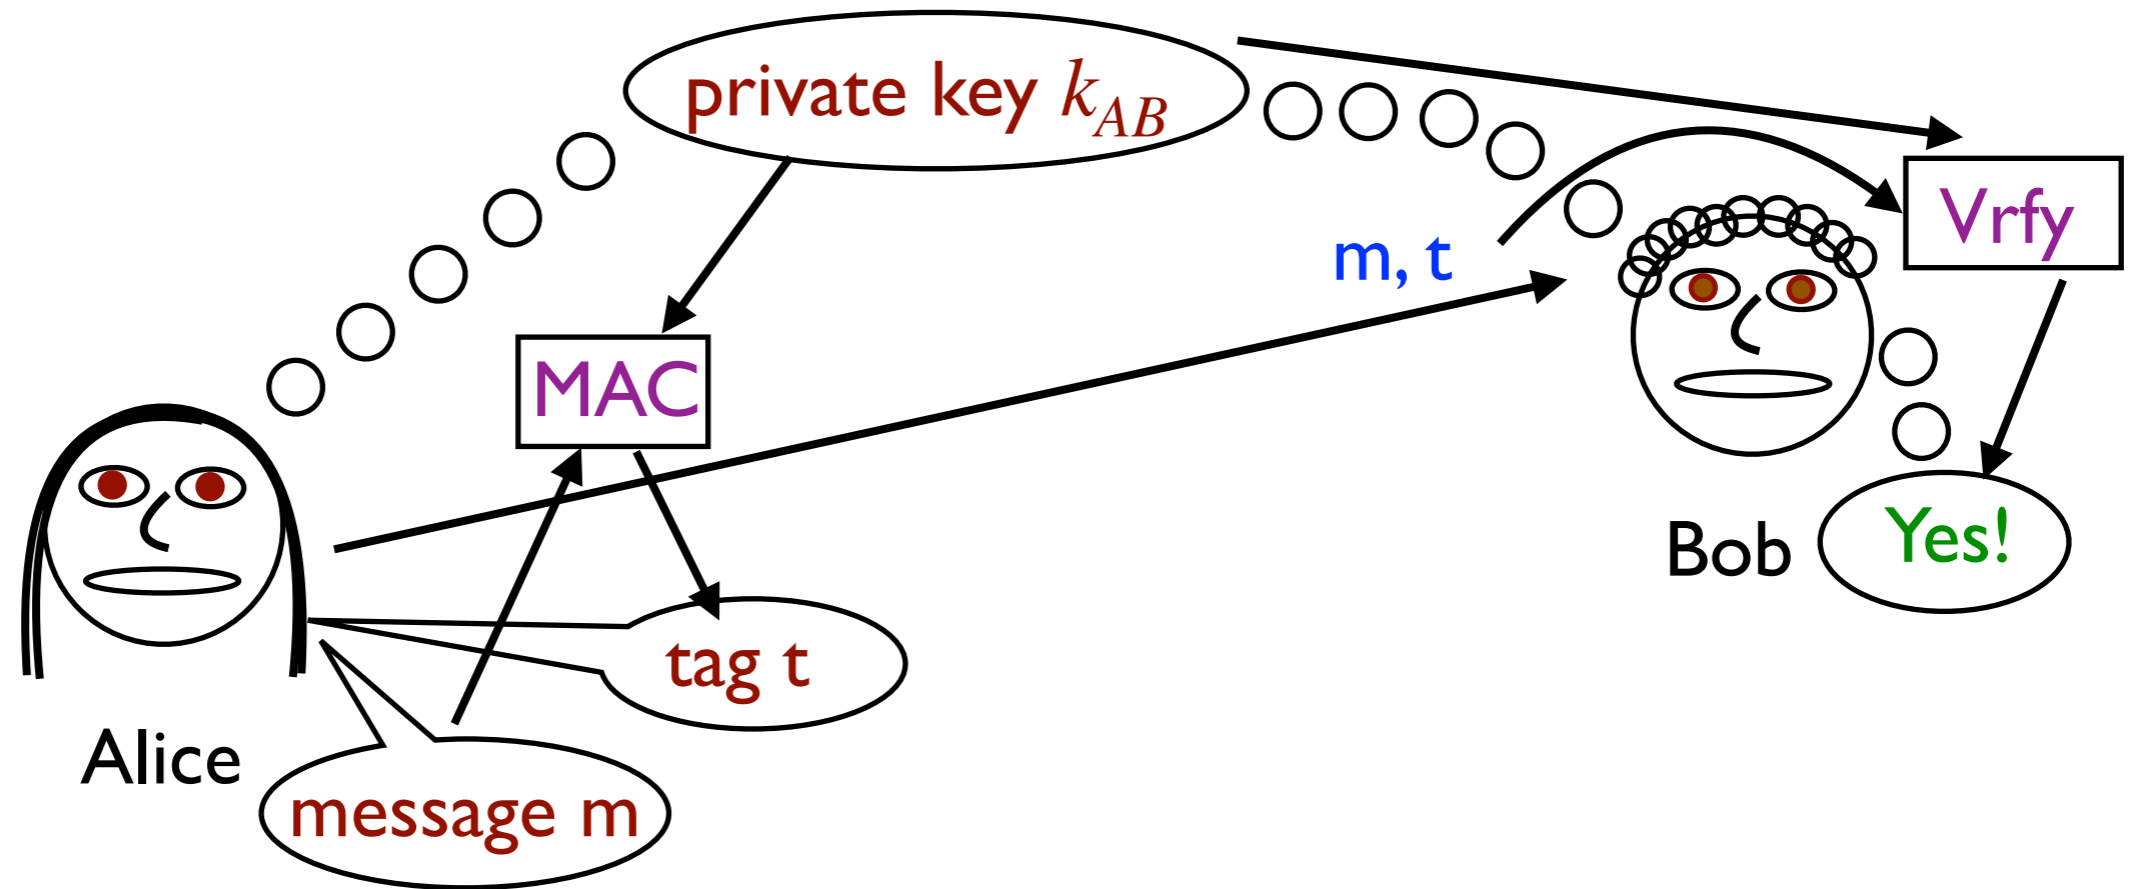
MACs Do Not Have Transferability



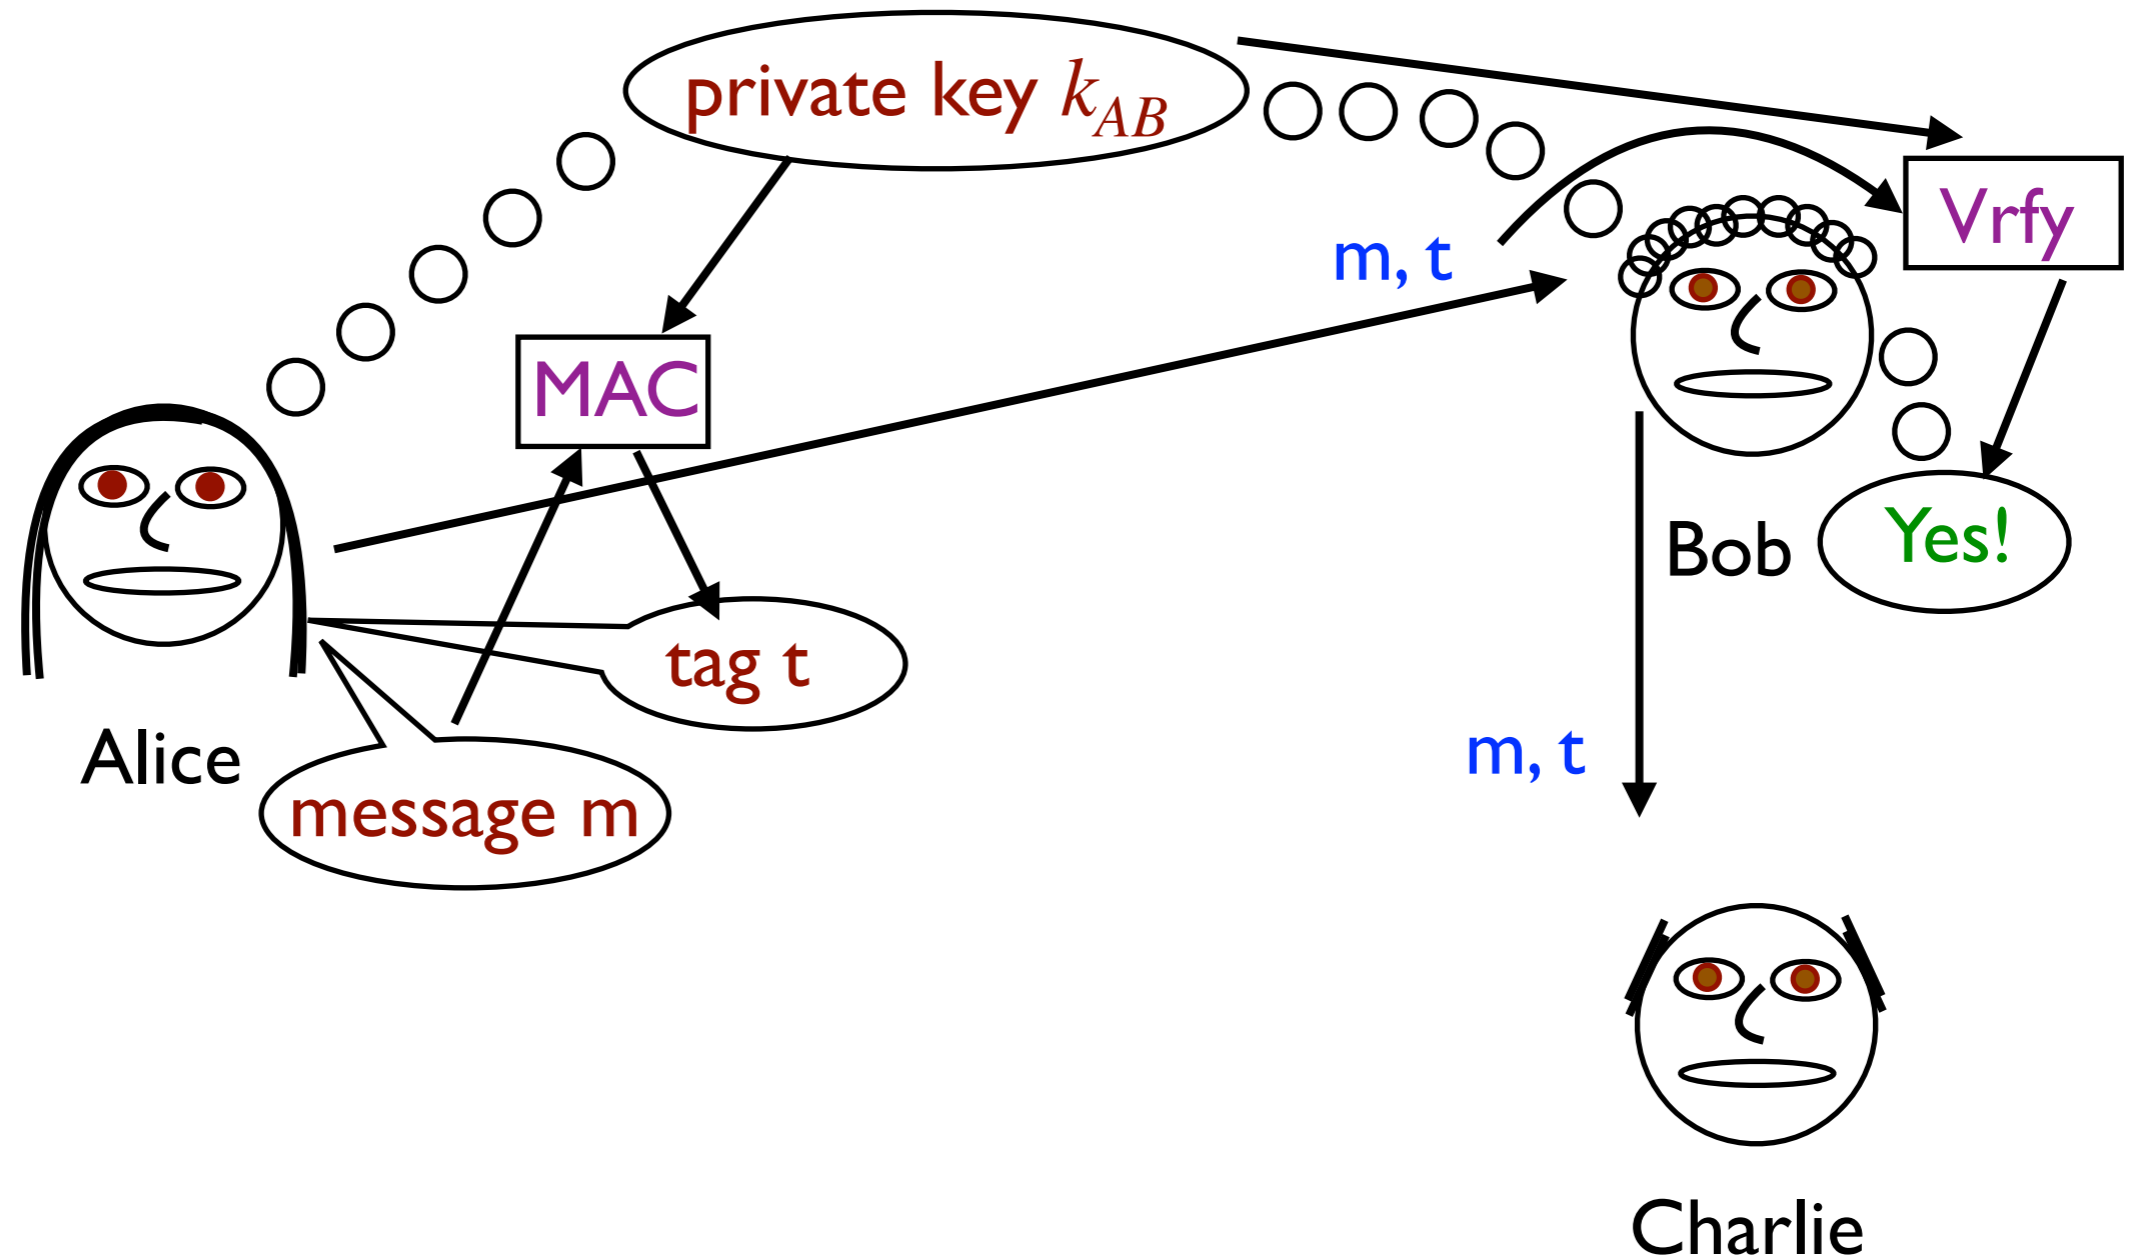
MACs Do Not Have Transferability



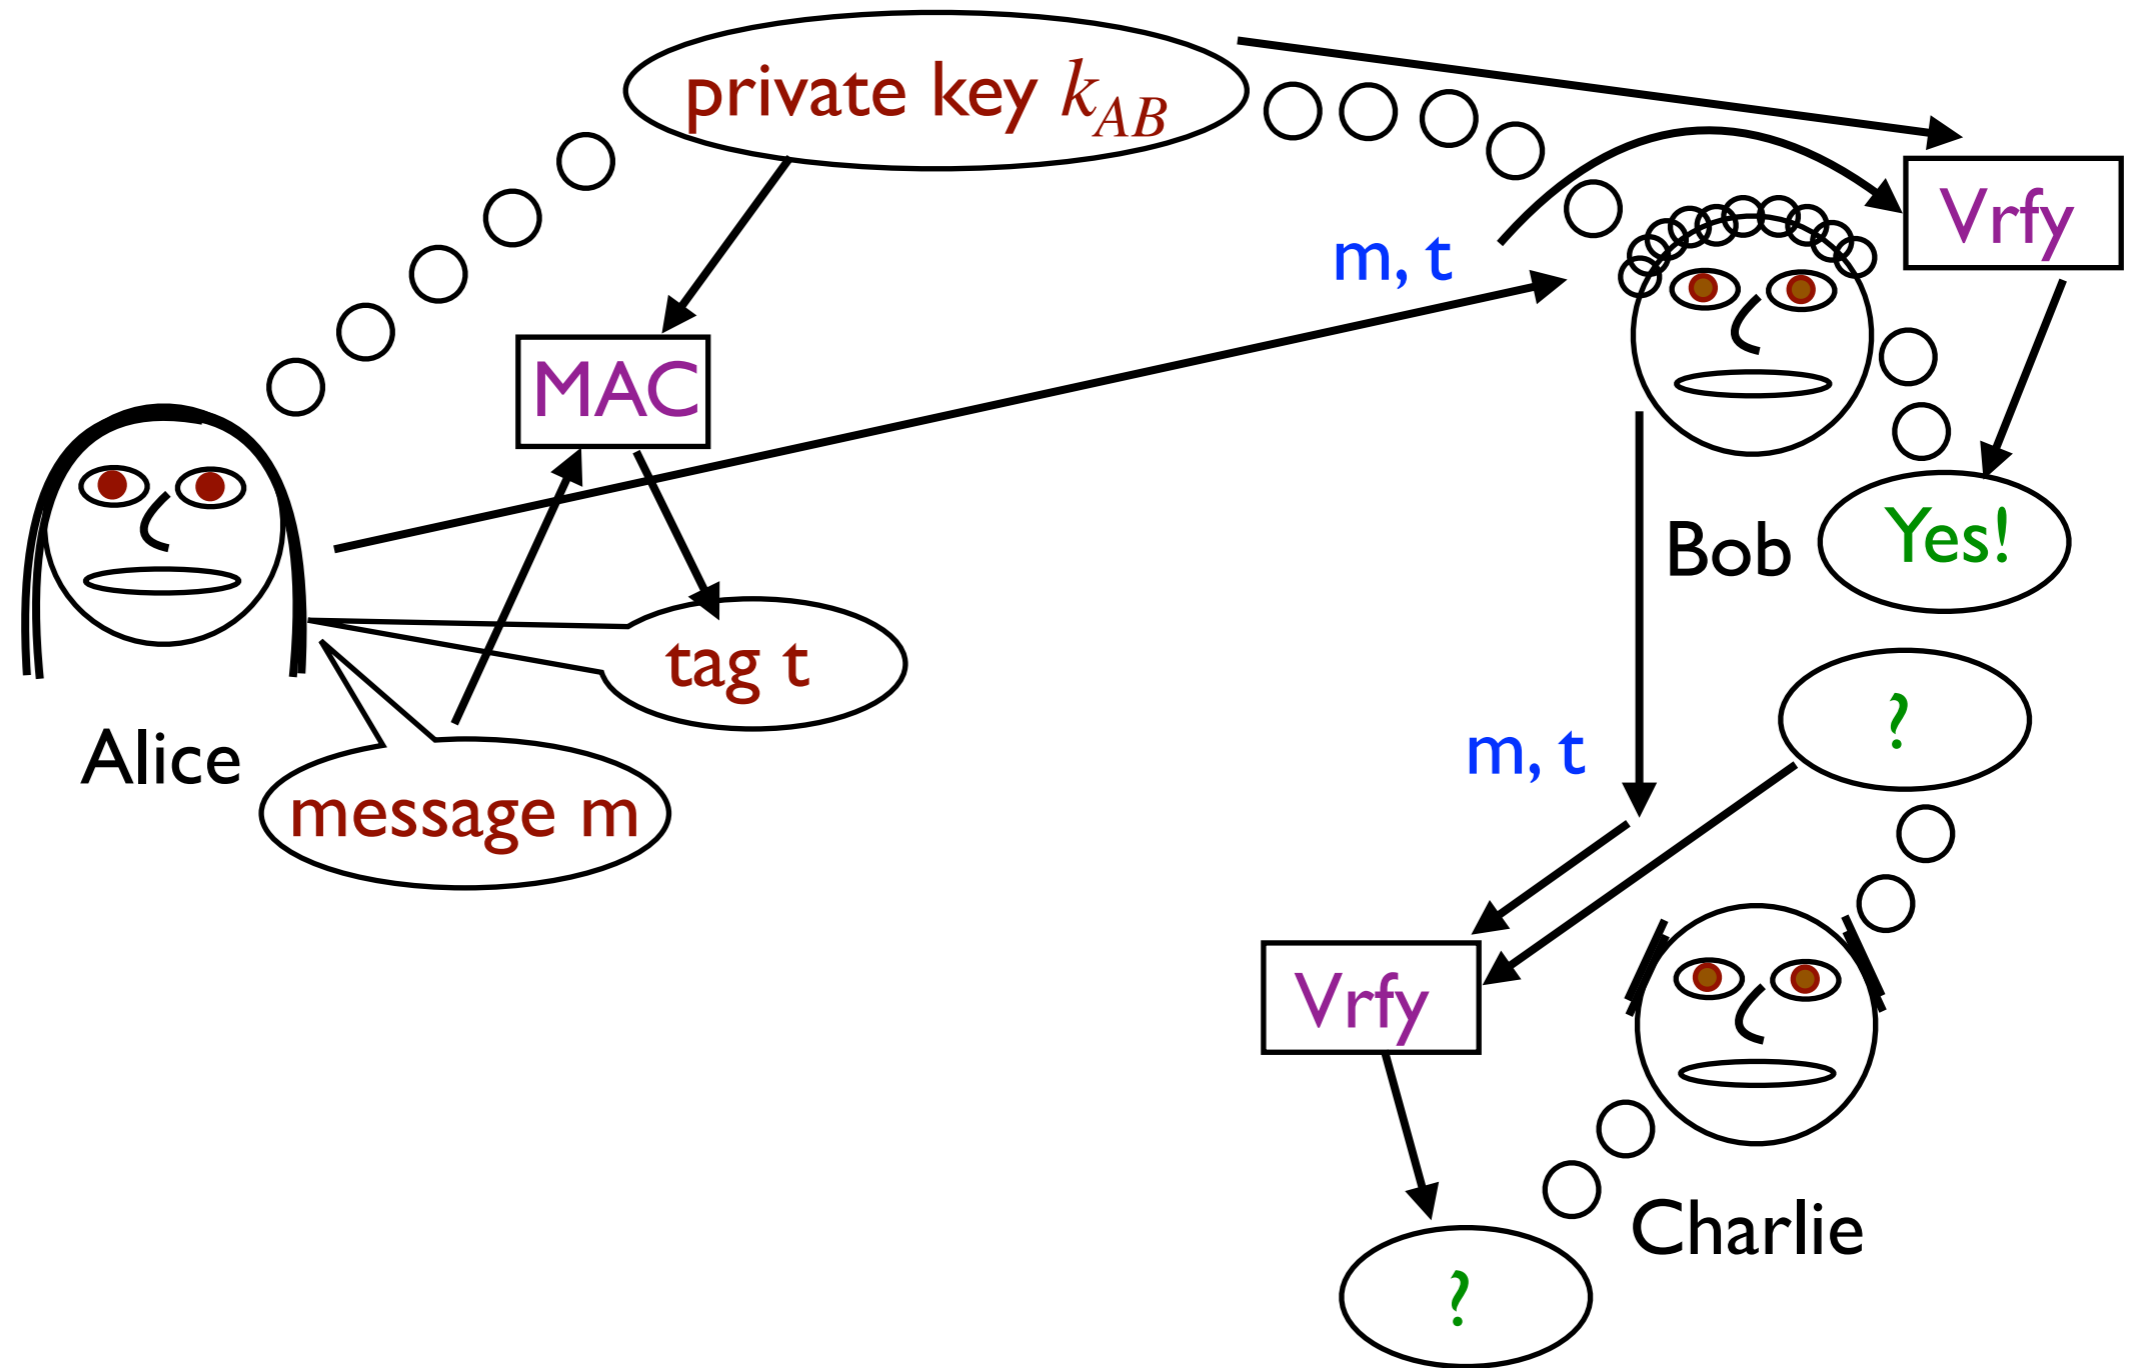
MACs Do Not Have Transferability



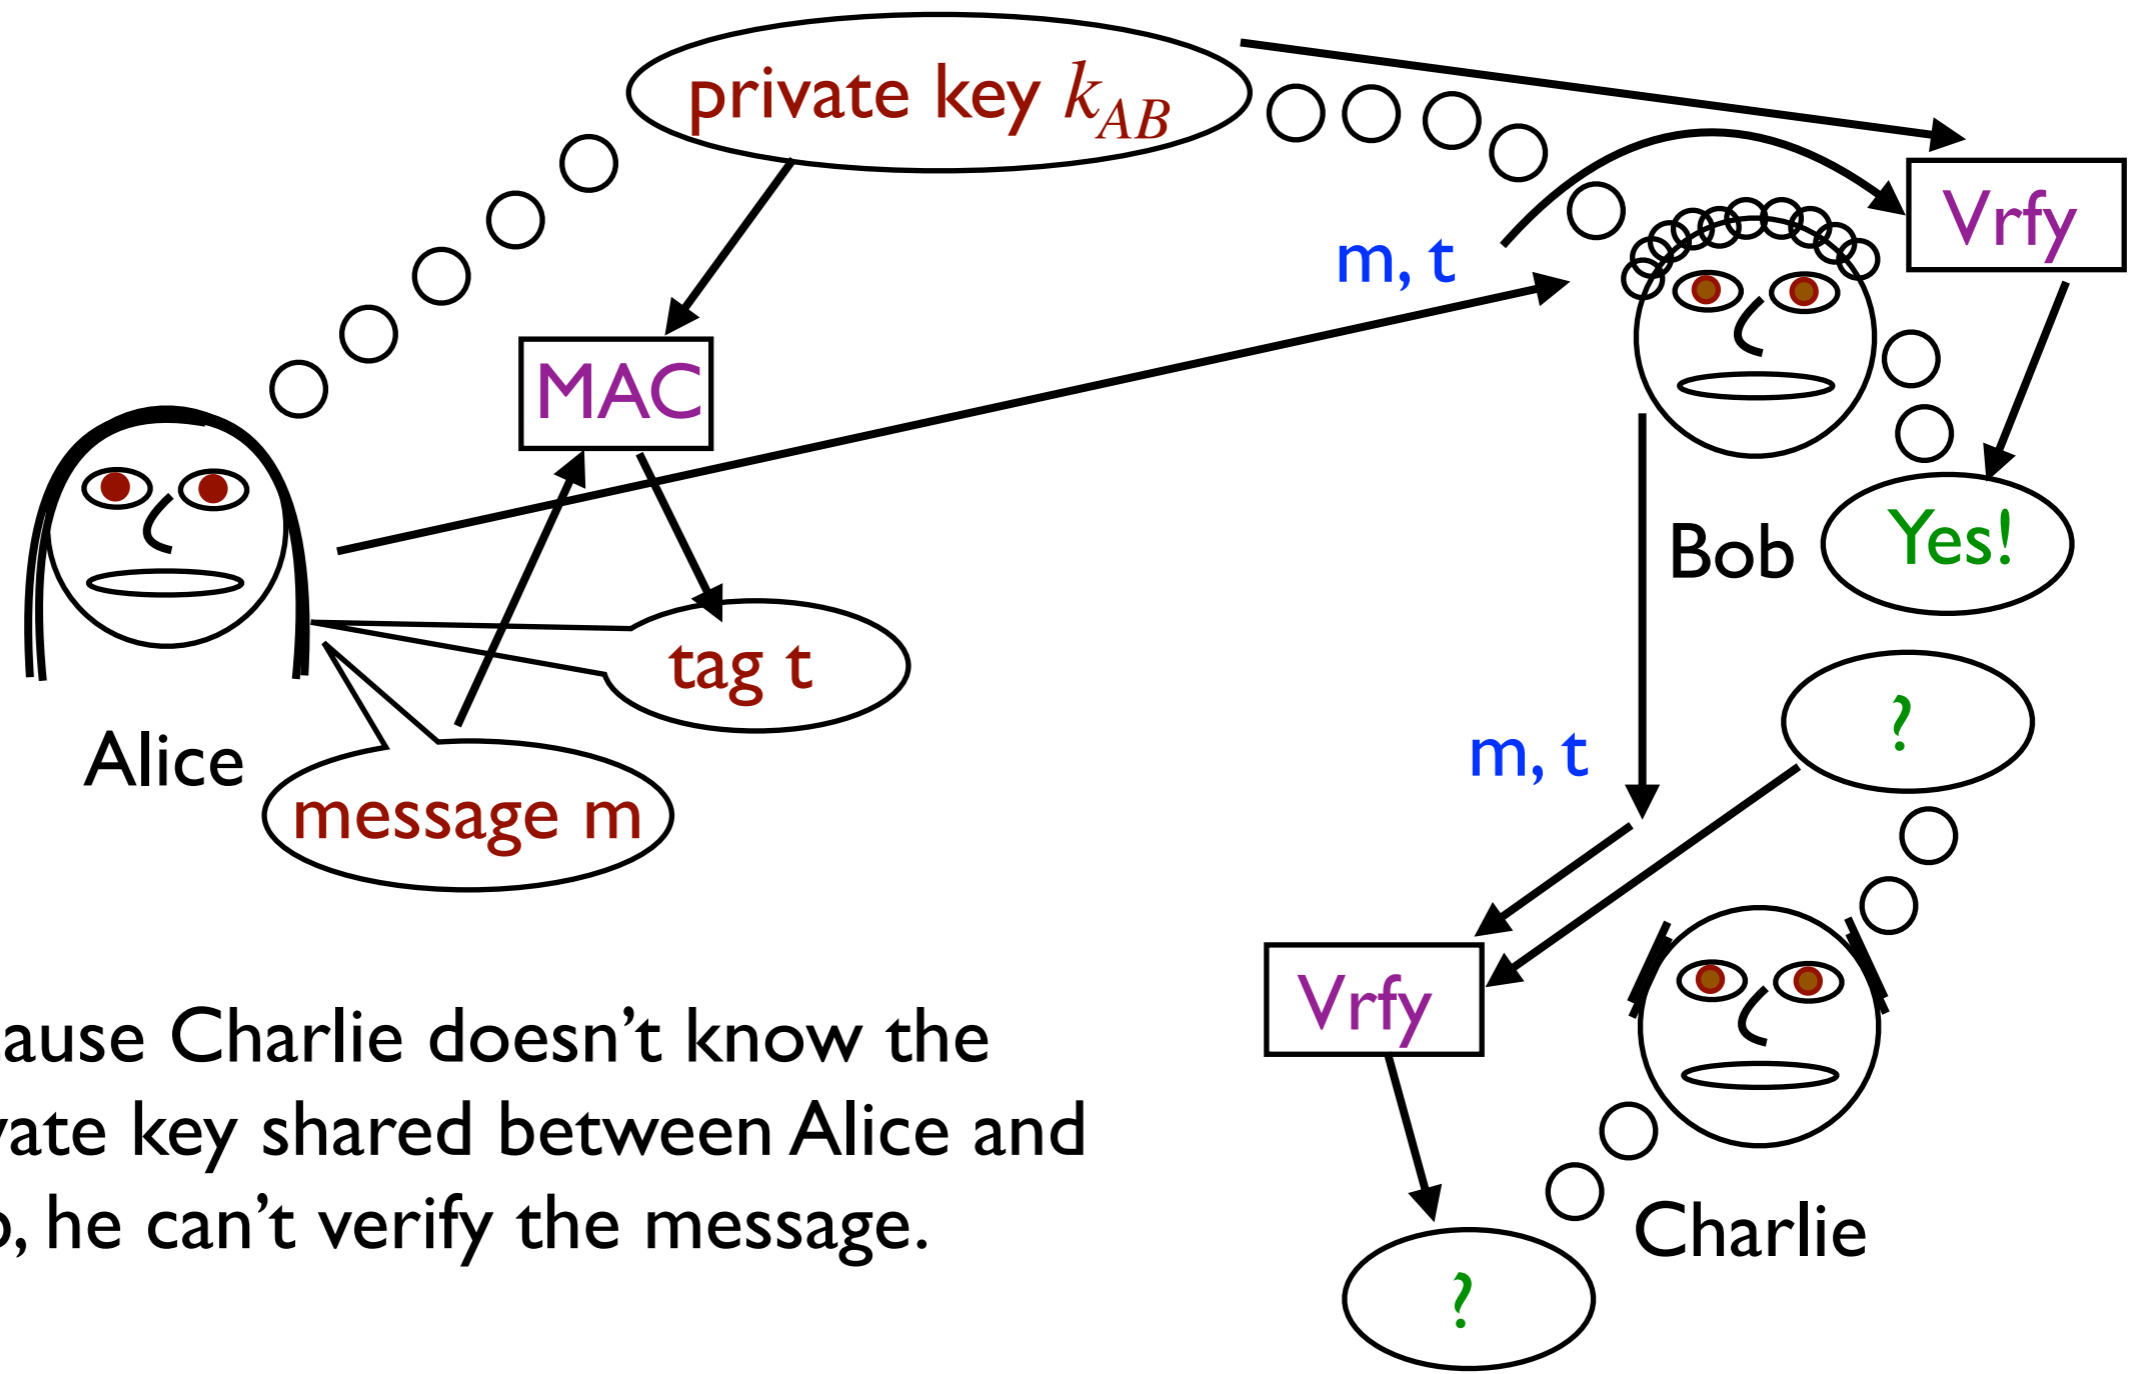
MACs Do Not Have Transferability



MACs Do Not Have Transferability




MACs Do Not Have Transferability



Because Charlie doesn't know the private key shared between Alice and Bob, he can't verify the message.


MACs Do Not Have Transferability

Suppose we do give Charlie a copy of the same private key shared by Alice and Bob.



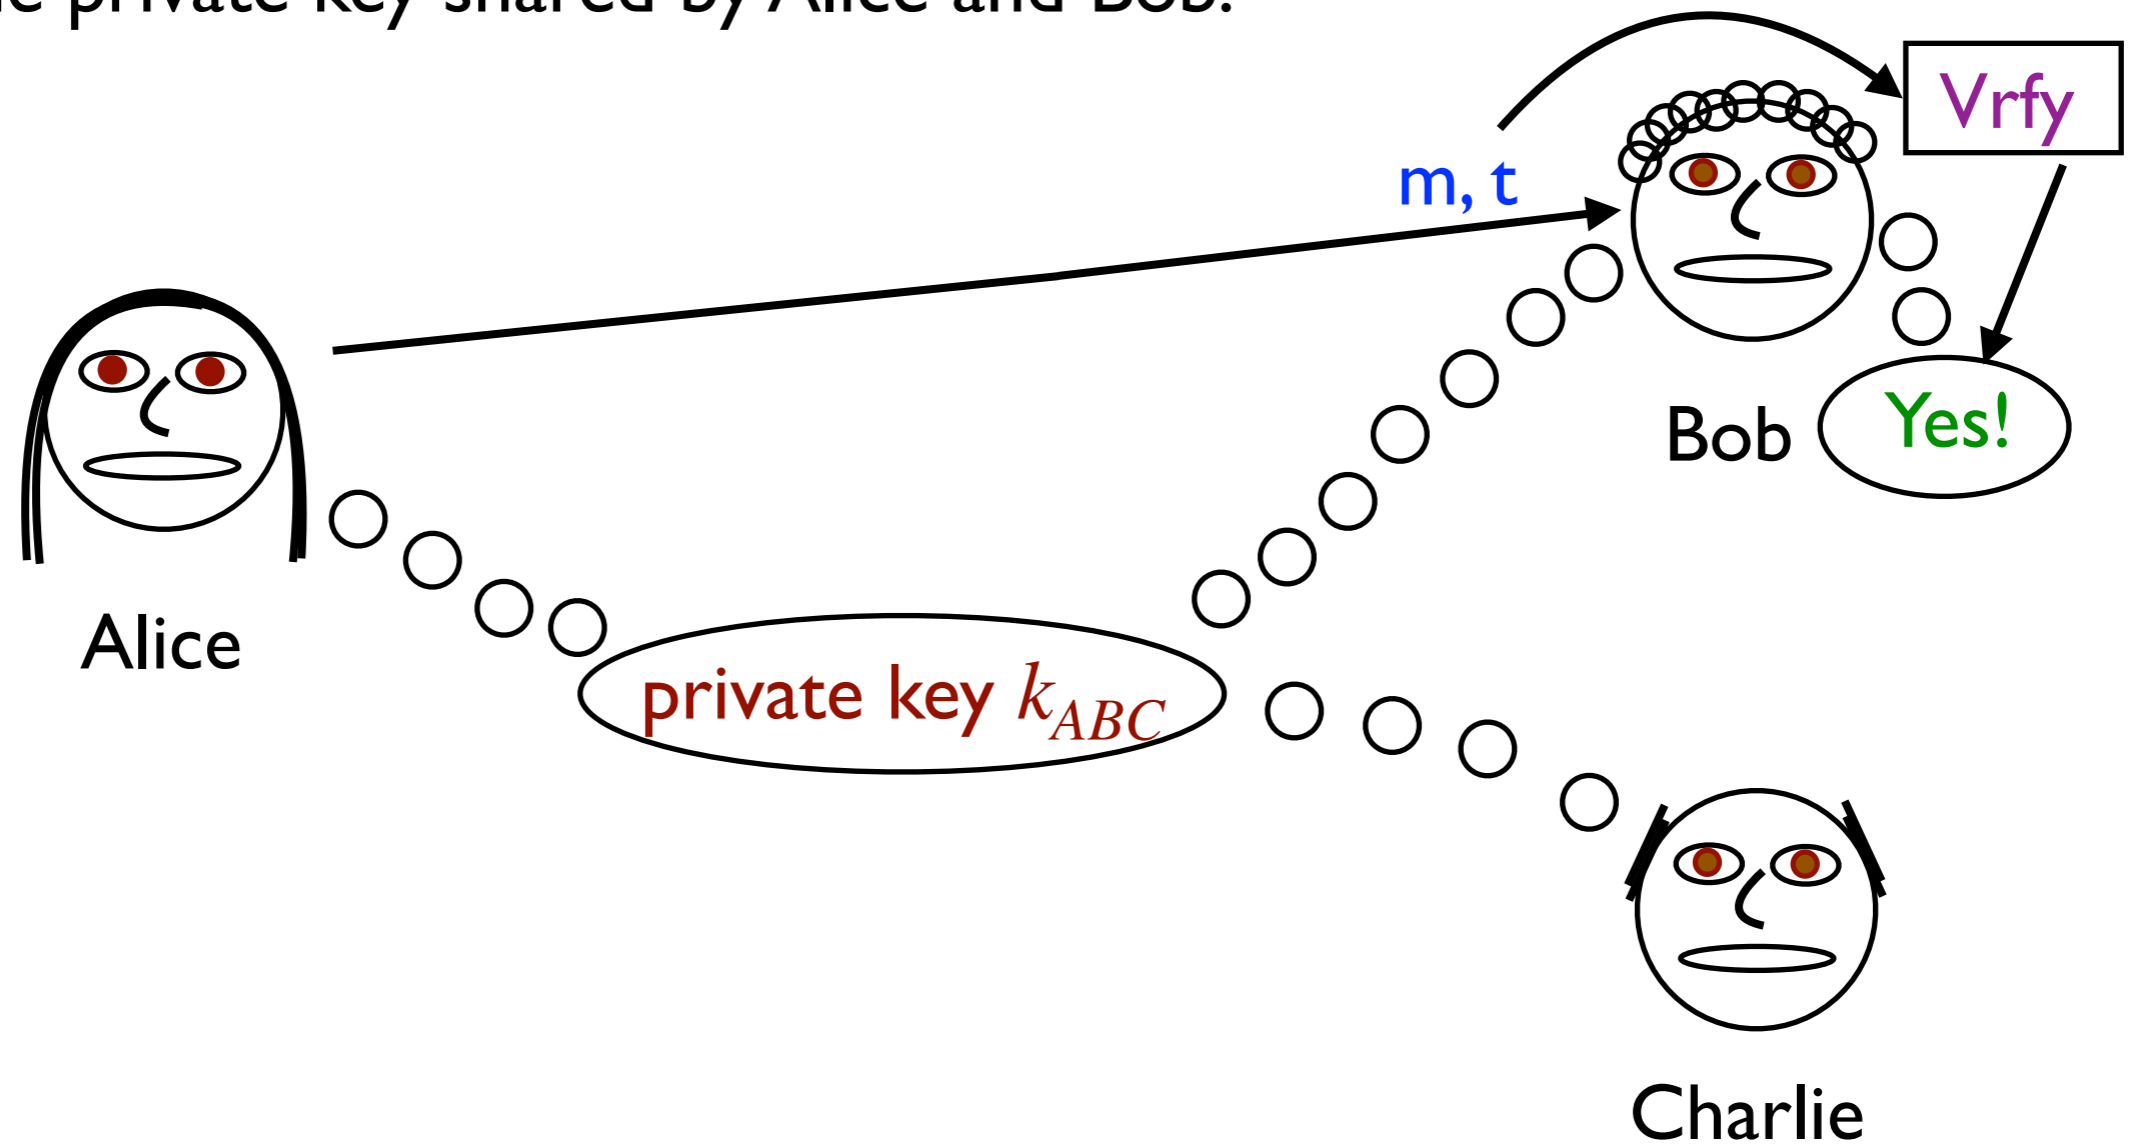
MACs Do Not Have Transferability

Suppose we do give Charlie a copy of the same private key shared by Alice and Bob.



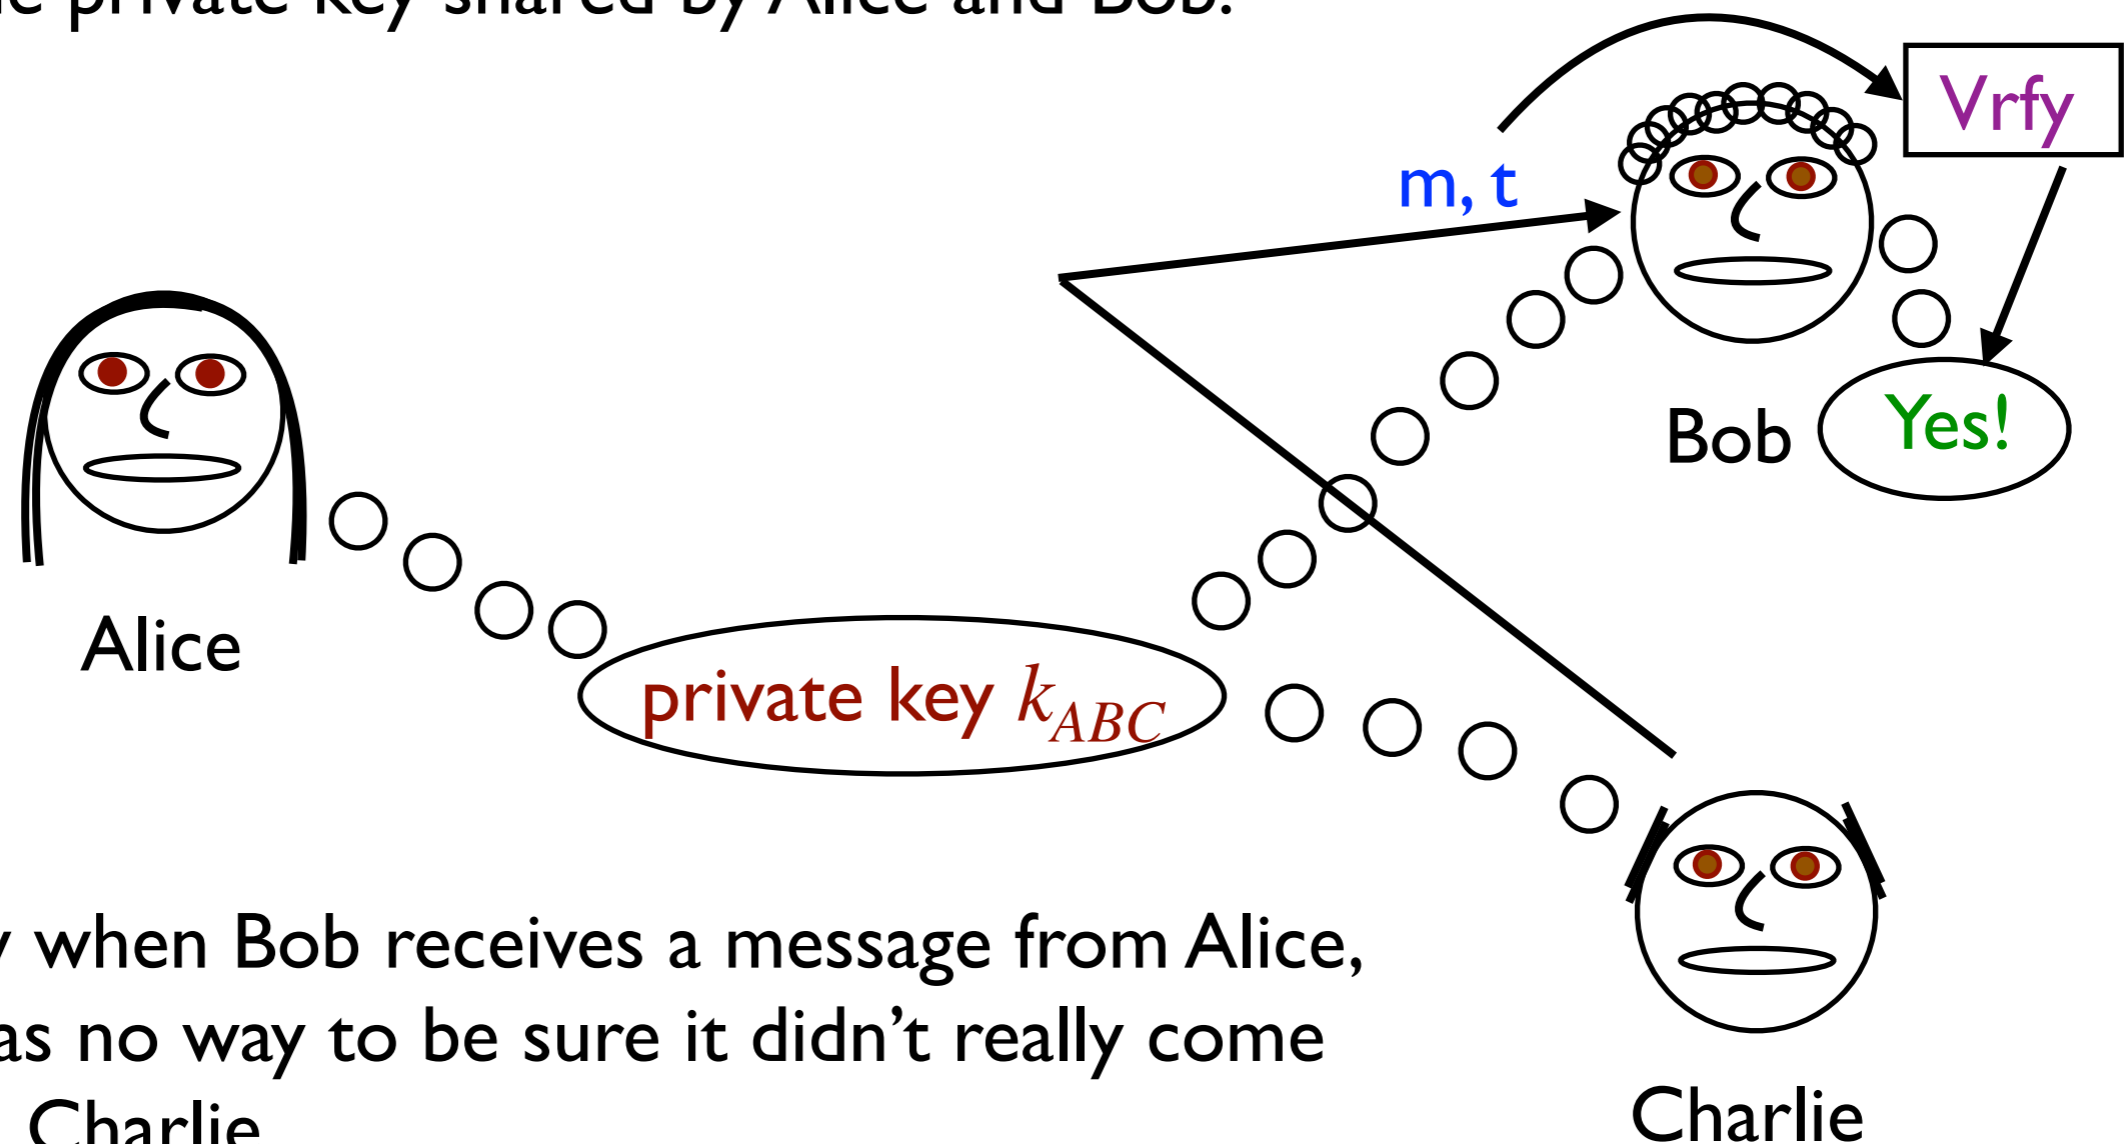
MACs Do Not Have Transferability

Suppose we do give Charlie a copy of the same private key shared by Alice and Bob.



MACs Do Not Have Transferability

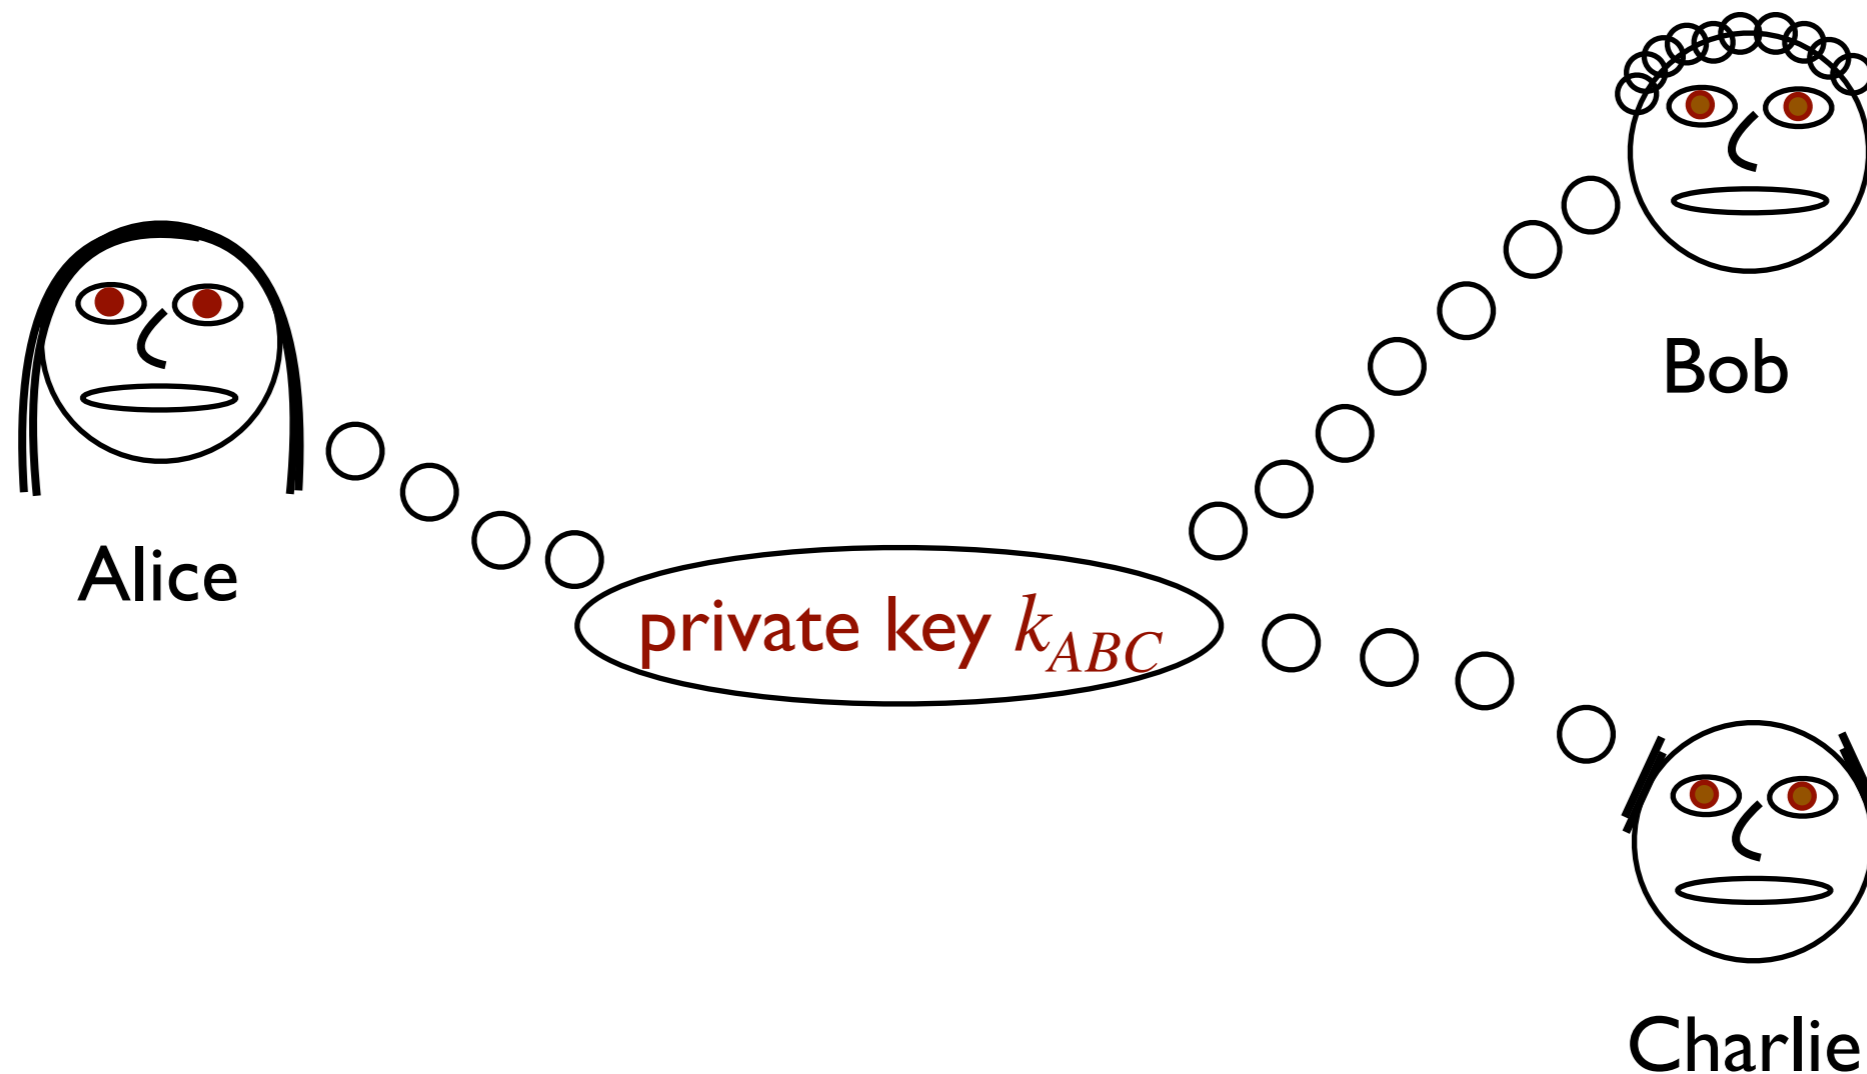
Suppose we do give Charlie a copy of the same private key shared by Alice and Bob.



Now when Bob receives a message from Alice, he has no way to be sure it didn't really come from Charlie.

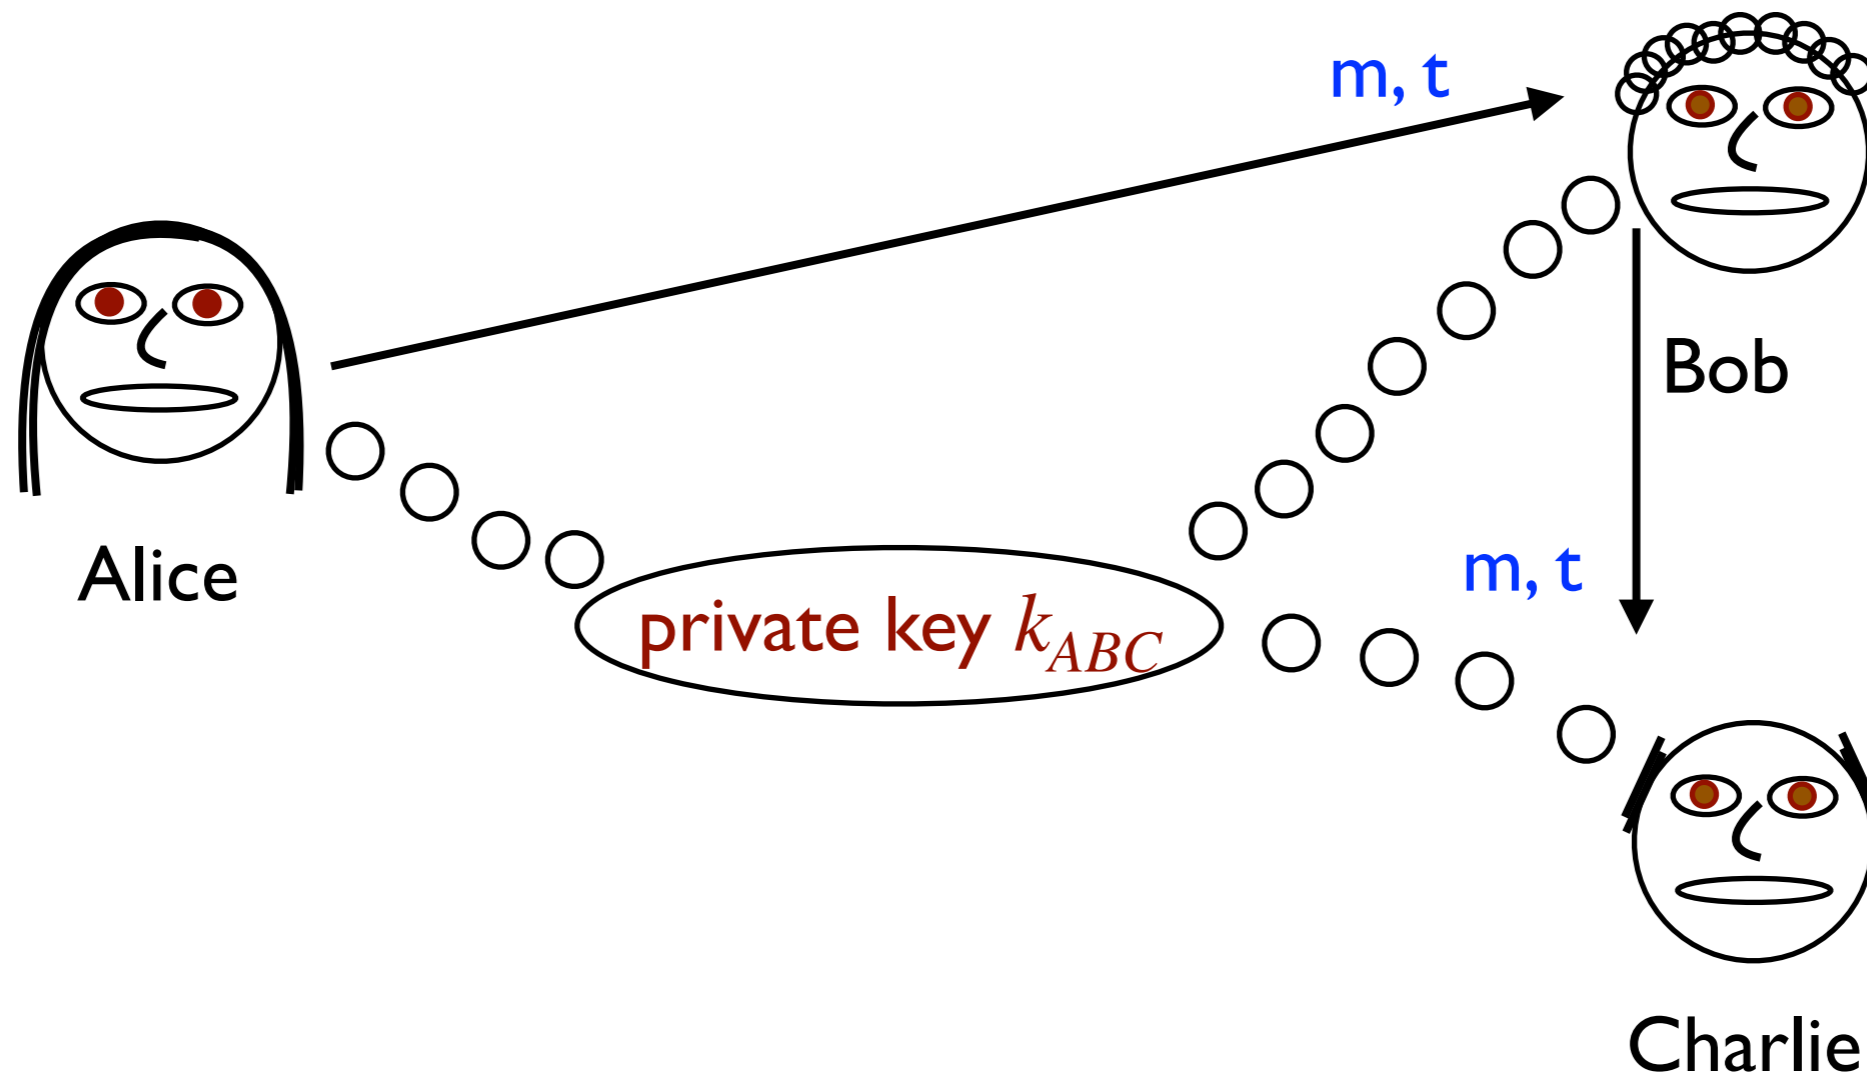
MACs Do Not Have Transferability

Suppose we do give Charlie a copy of the same private key shared by Alice and Bob.



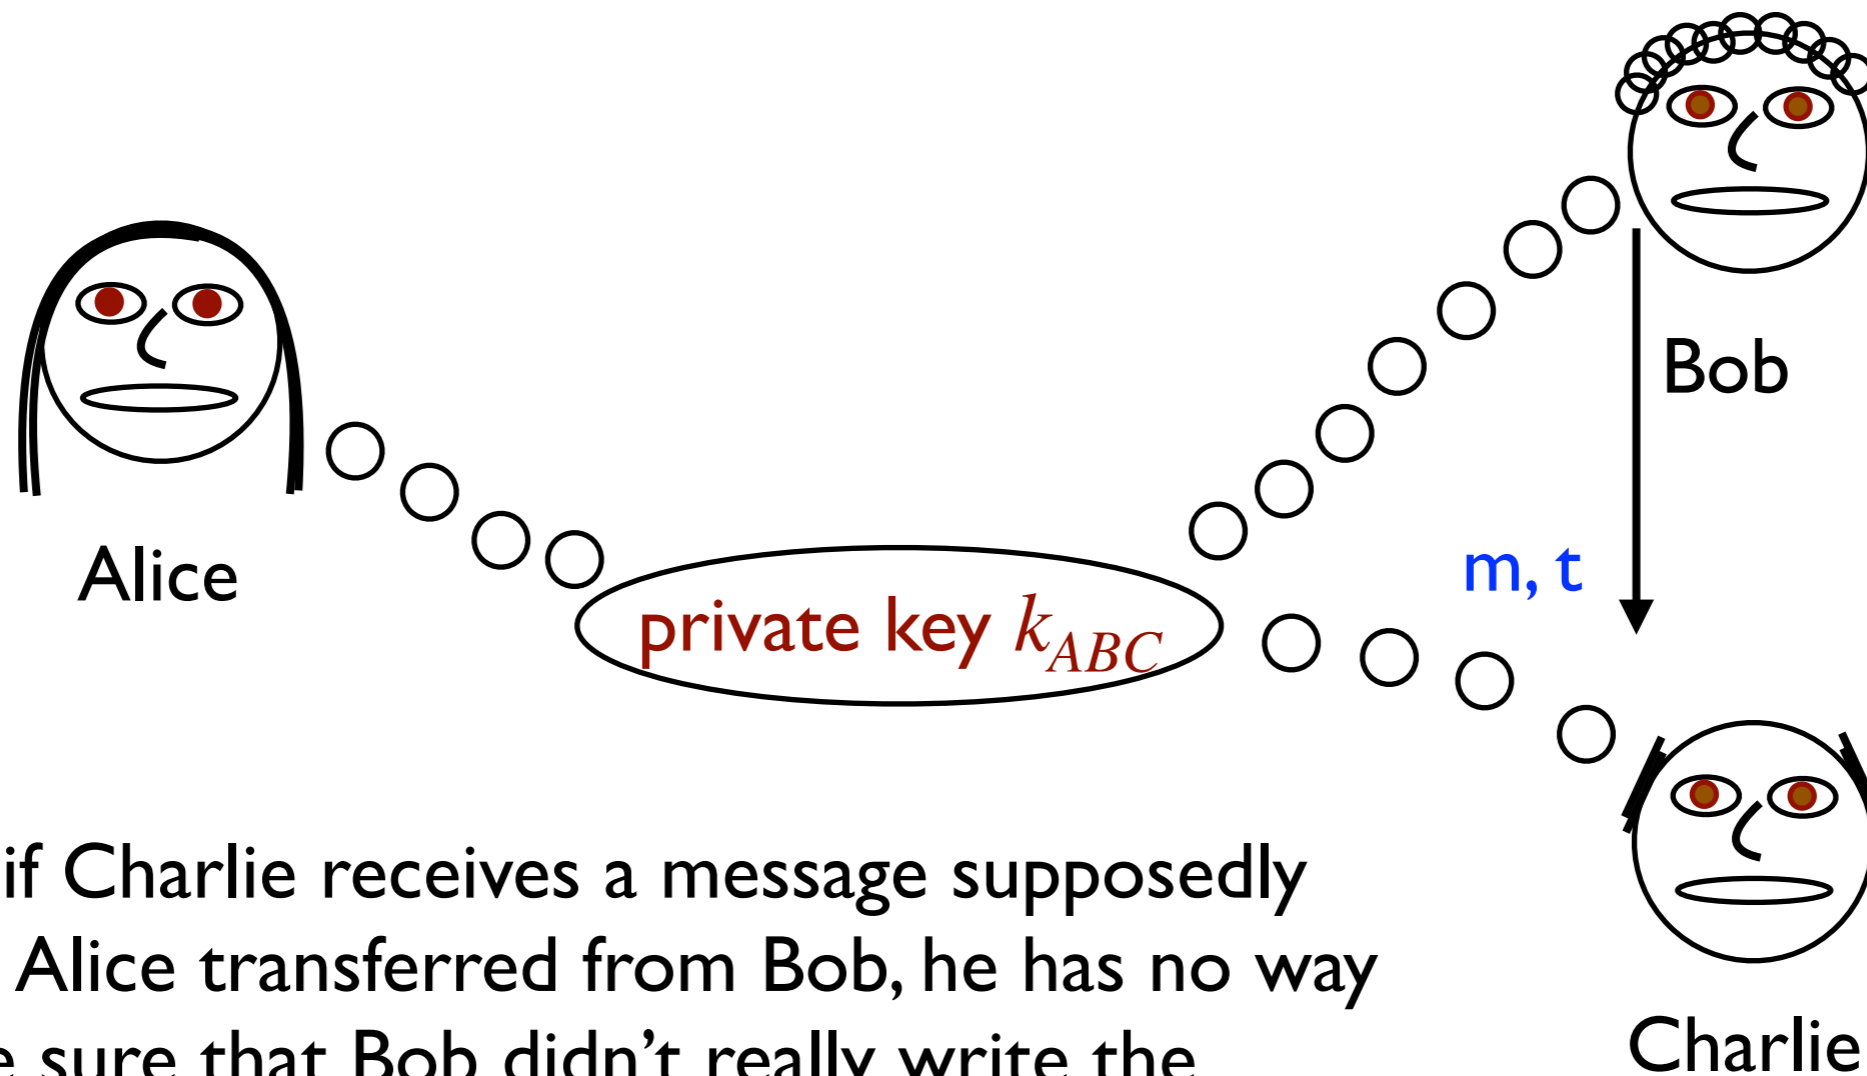
MACs Do Not Have Transferability

Suppose we do give Charlie a copy of the same private key shared by Alice and Bob.



MACs Do Not Have Transferability

Suppose we do give Charlie a copy of the same private key shared by Alice and Bob.



And if Charlie receives a message supposedly from Alice transferred from Bob, he has no way to be sure that Bob didn't really write the message himself.

Non-Repudiation

The **non-repudiation** property of digital signatures says that Alice cannot later deny having sent the message.

In particular, Bob can present the signed message to a third party Charlie who also has Alice's public key and Charlie will agree with Bob that the message is valid and was signed by Alice.

This is important, e.g., for legal contracts:

If necessary, Bob can prove in court that Alice agreed to the contract.

Non-Repudiation

The **non-repudiation** property of digital signatures says that Alice cannot later deny having sent the message.

In particular, Bob can present the signed message to a third party Charlie who also has Alice's public key and Charlie will agree with Bob that the message is valid and was signed by Alice.

This is important, e.g., for legal contracts:

If necessary, Bob can prove in court that Alice agreed to the contract.

Note that a MAC does not have this property: Anyone who can verify can forge messages.

Non-Repudiation

The **non-repudiation** property of digital signatures says that Alice cannot later deny having sent the message.

In particular, Bob can present the signed message to a third party Charlie who also has Alice's public key and Charlie will agree with Bob that the message is valid and was signed by Alice.

This is important, e.g., for legal contracts:

If necessary, Bob can prove in court that Alice agreed to the contract.

Note that a MAC does not have this property: Anyone who can verify can forge messages.

Discuss: Are digital signatures more or less secure than hardcopy signatures? (More secure/Less secure/Similar security)

