



# TumbleBit: An Untrusted Bitcoin-Compatible Anonymous Payment Hub

Ethan Heilman, Leen AlShenibr, Foteini Baldimtsi,  
Alessandra Scafuro, Sharon Goldberg



Scaling Bitcoin Milan 2016



# Introduction

## TumbleBit is:

1. **Private:** Unlinkable Bitcoin payments and k-anonymous mixing,
2. **Untrusted:** No one including Tumbler can steal or link payments.
3. **Scalable (payment hub):** scales transaction velocity and volume.
4. **Compatible:** Works with today's Bitcoin protocol.

## Why is compatibility hard?

Our protocol must work with highly constrained Bitcoin scripts which provide two very limited cryptographic operations.

## Two ways to use TumbleBit:

### TumbleBit can be used as a classic Bitcoin tumbler:

- k-anonymity within a mix,
- 4 transactions confirmed in 2 blocks (~20mins)

### When TumbleBit is used as a payment hub:

- Unlinkability within the payment phase,
- Payments confirmed in seconds,
- Payments are off-blockchain,  
... don't take up space on the blockchain.

# Introduction

## TumbleBit is:

1. **Private:** Unlinkable Bitcoin payments and k-anonymous mixing,
2. **Untrusted:** No one including Tumbler can steal or link payments.
3. **Scalable (payment hub):** scales transaction velocity and volume.
4. **Compatible:** Works with today's Bitcoin protocol.

## Why is compatibility hard?

Our protocol must work with highly constrained Bitcoin scripts which provide two very limited cryptographic operations.

## Two ways to use TumbleBit:



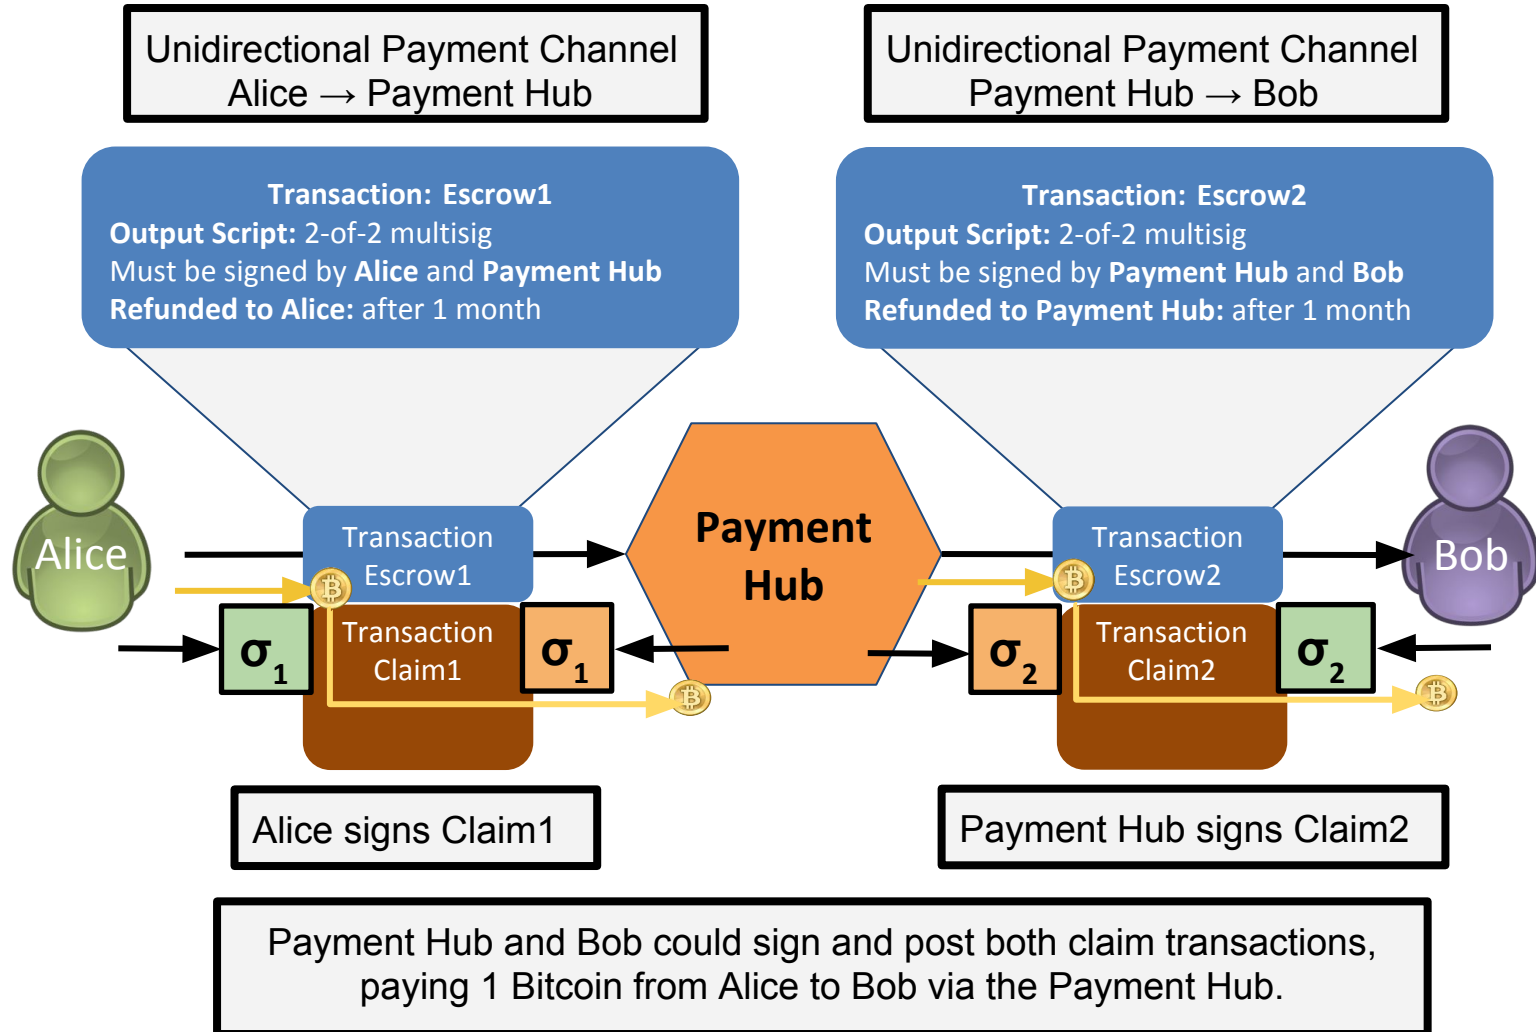
When used as a payment hub, TumbleBit helps scale Bitcoin's transaction velocity (faster payments), and transaction volume (higher maximum payments).

### When TumbleBit is used as a payment hub:

- Unlinkability within the payment phase,
- Payments confirmed in seconds,
- Payments are off-blockchain,  
... don't take up space on the blockchain.

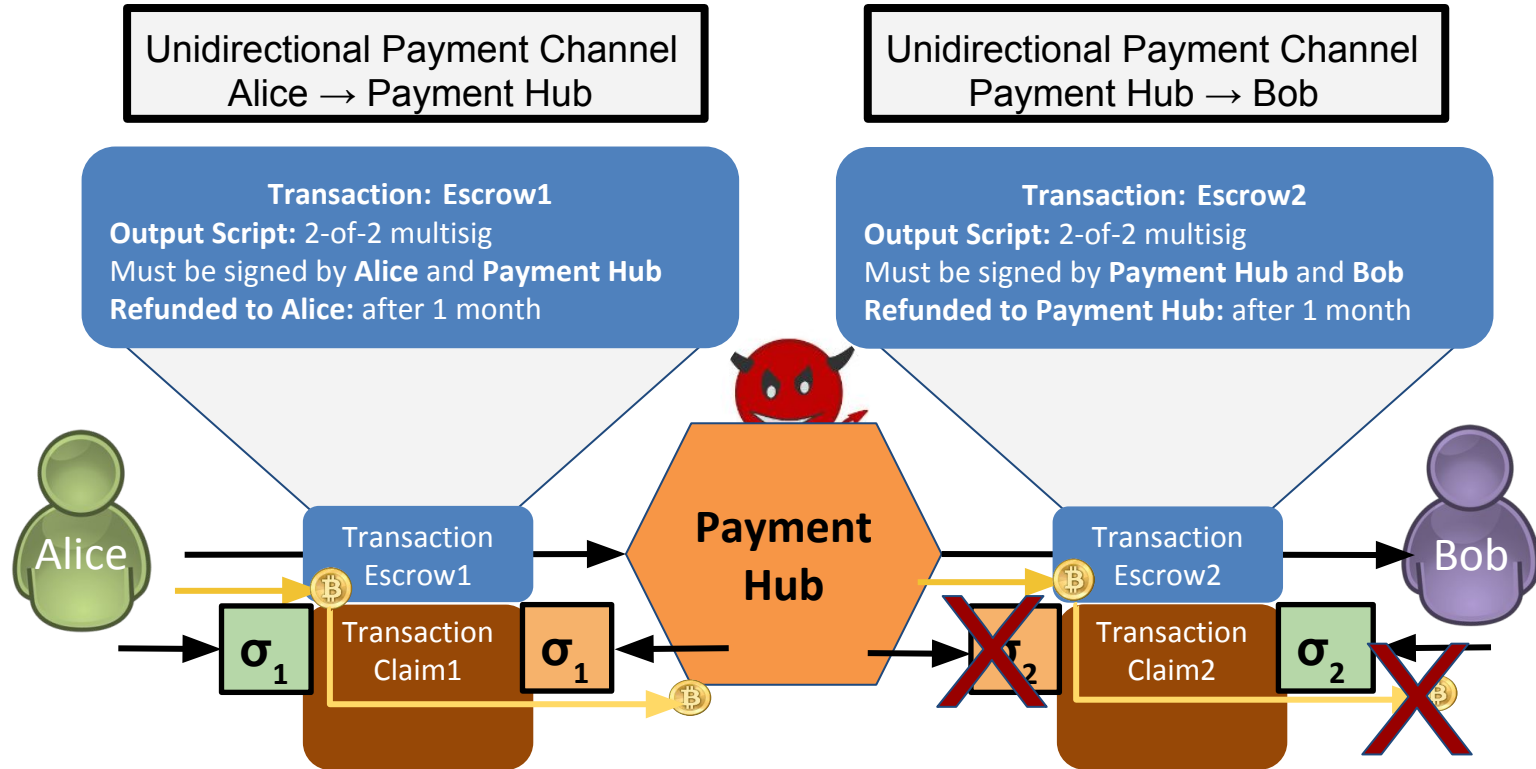
# Background: Payment Hub

**A payment hub:** routes payment channels.



# Background: Payment Hub

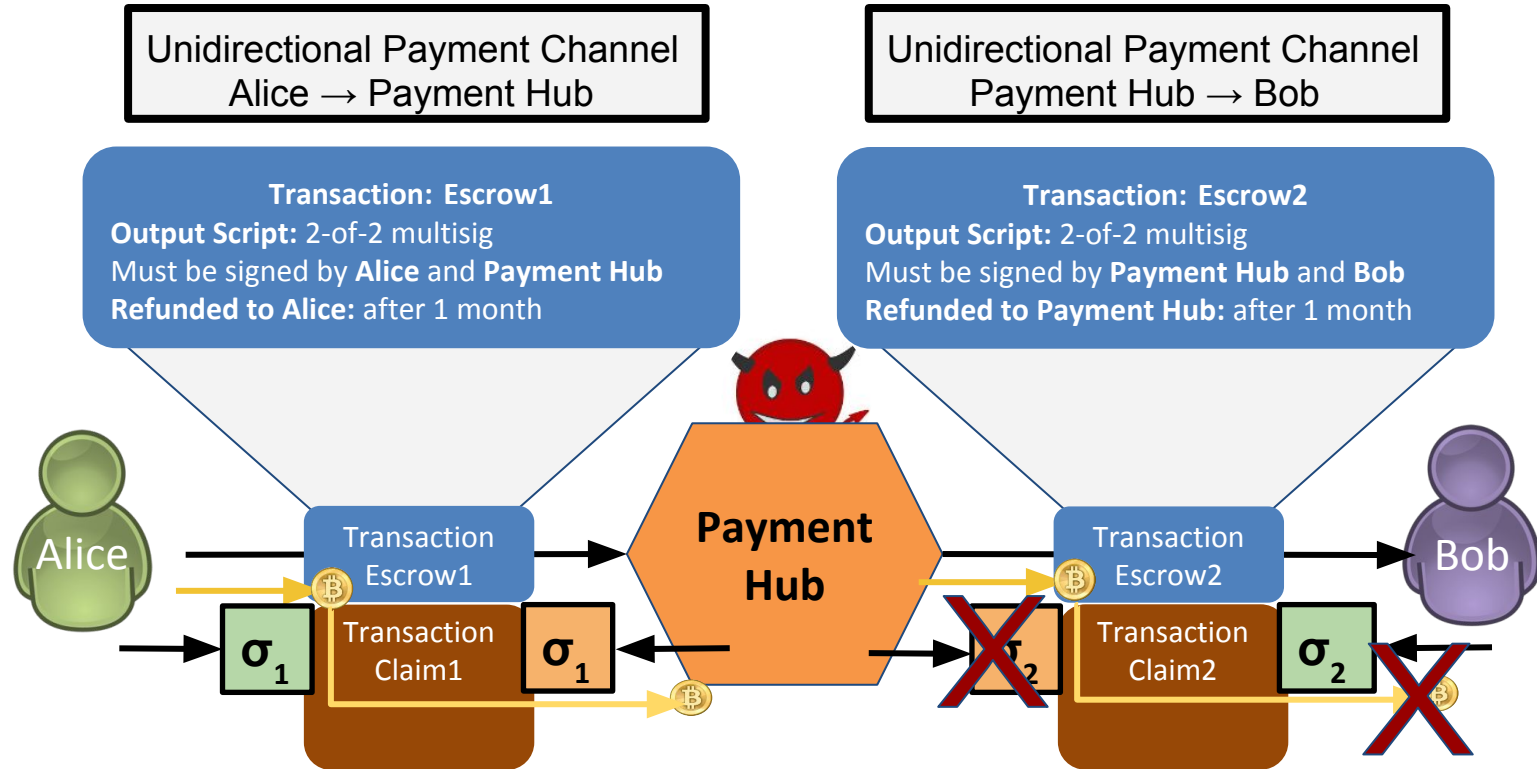
**A payment hub:** routes payment channels.



**...But what if the hub is malicious,  
and takes Alice's bitcoin and doesn't pay Bob?**

# Background: Payment Hub

**A payment hub:** routes payment channels.



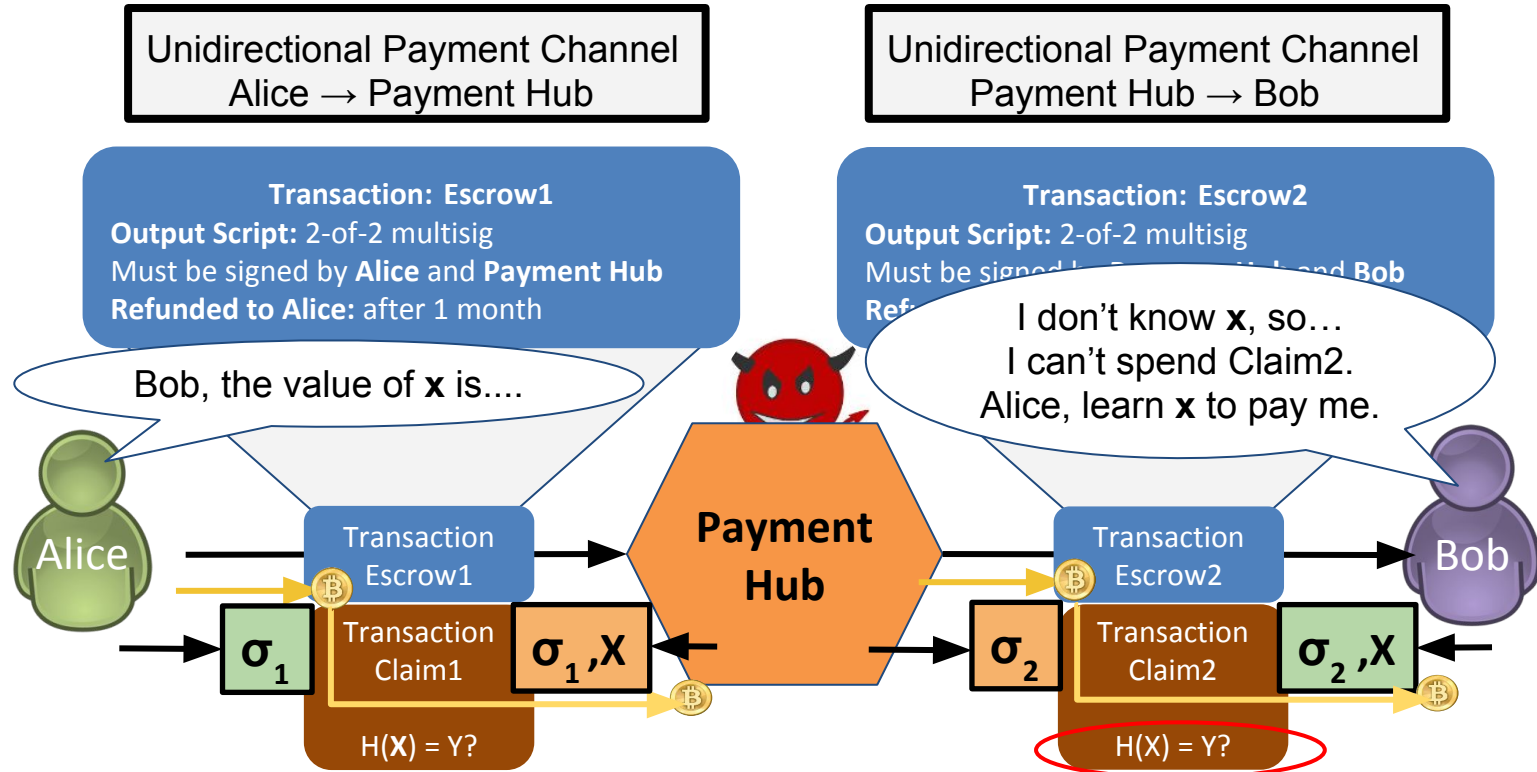
**...But what if the hub is malicious,**

**Atomicity:** If Claim1 and Claim2 happen atomically then theft is prevented.

Hash locks provide this property.

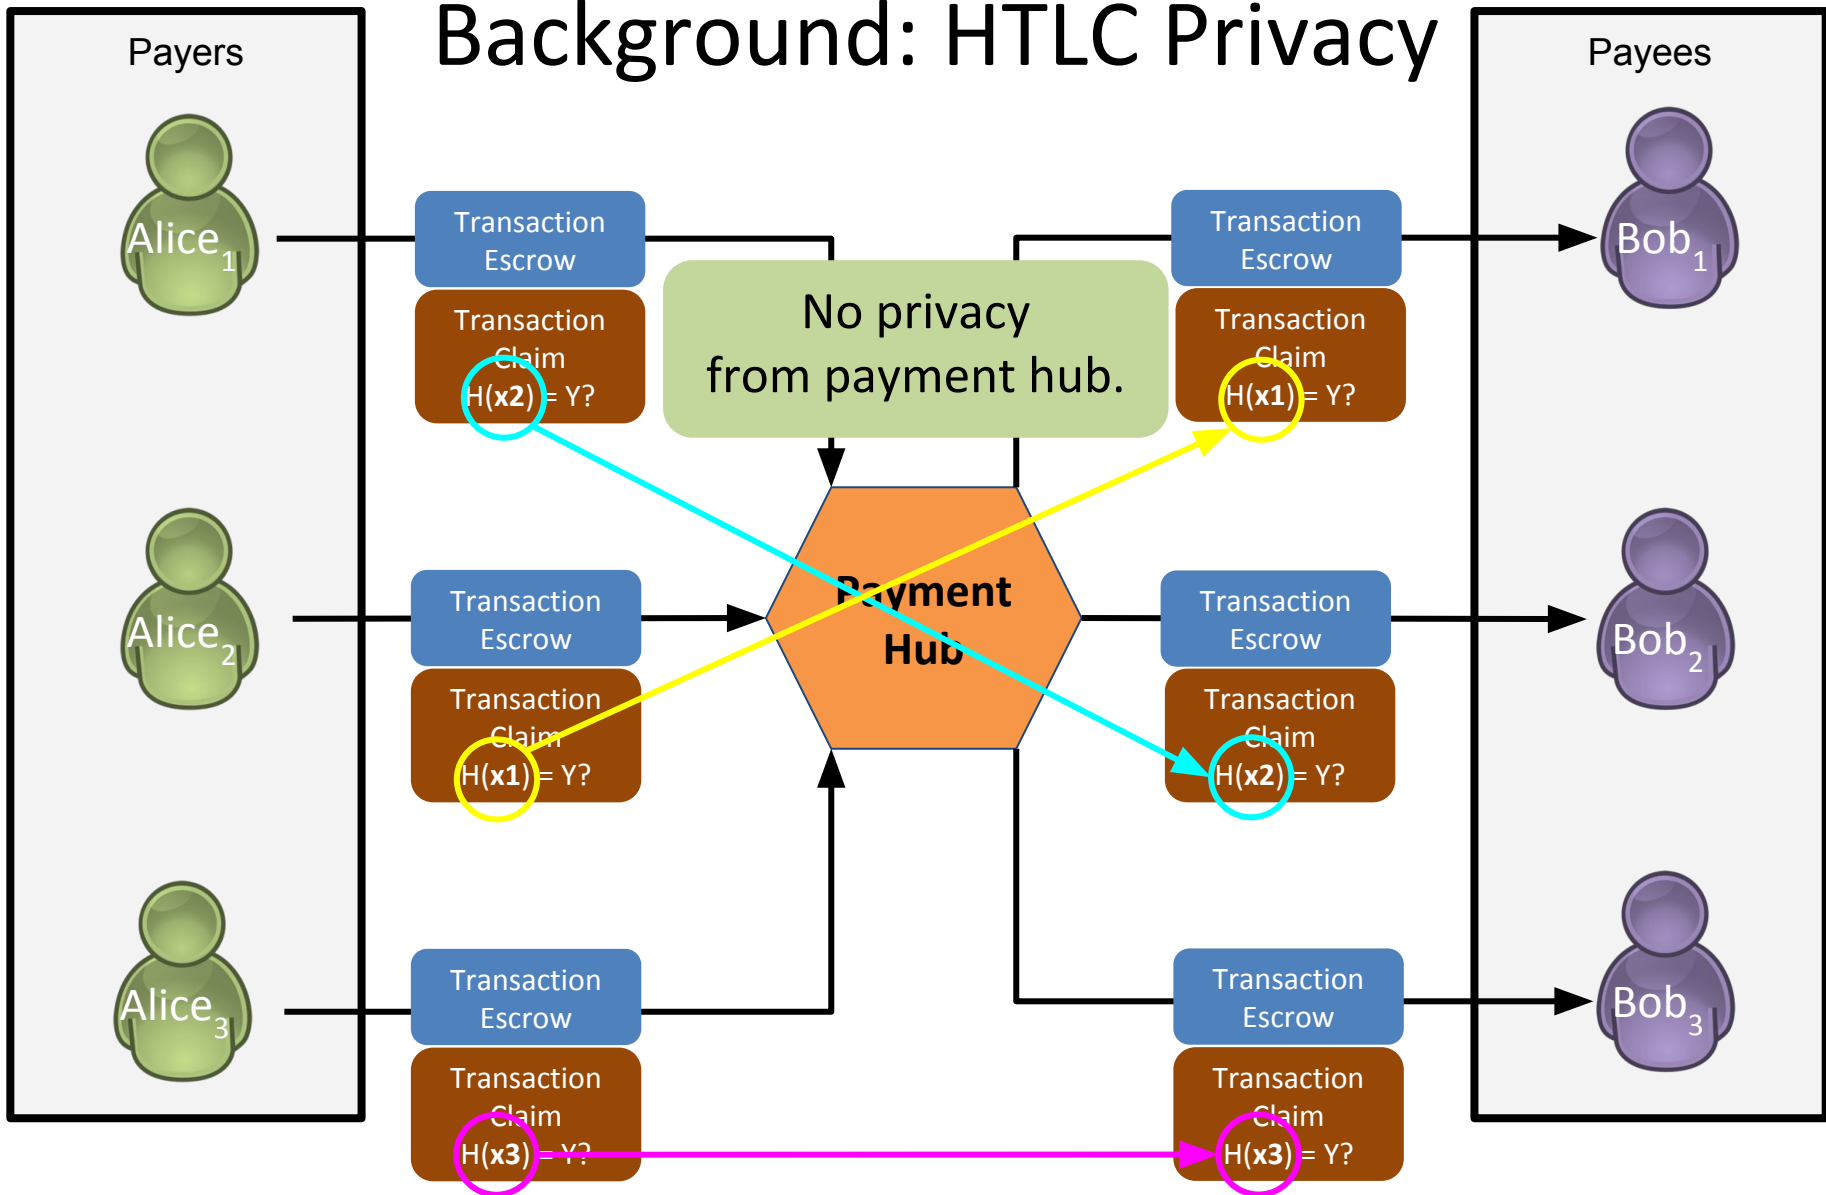
# Background: Payment Hub

A payment hub: routes payment channels.



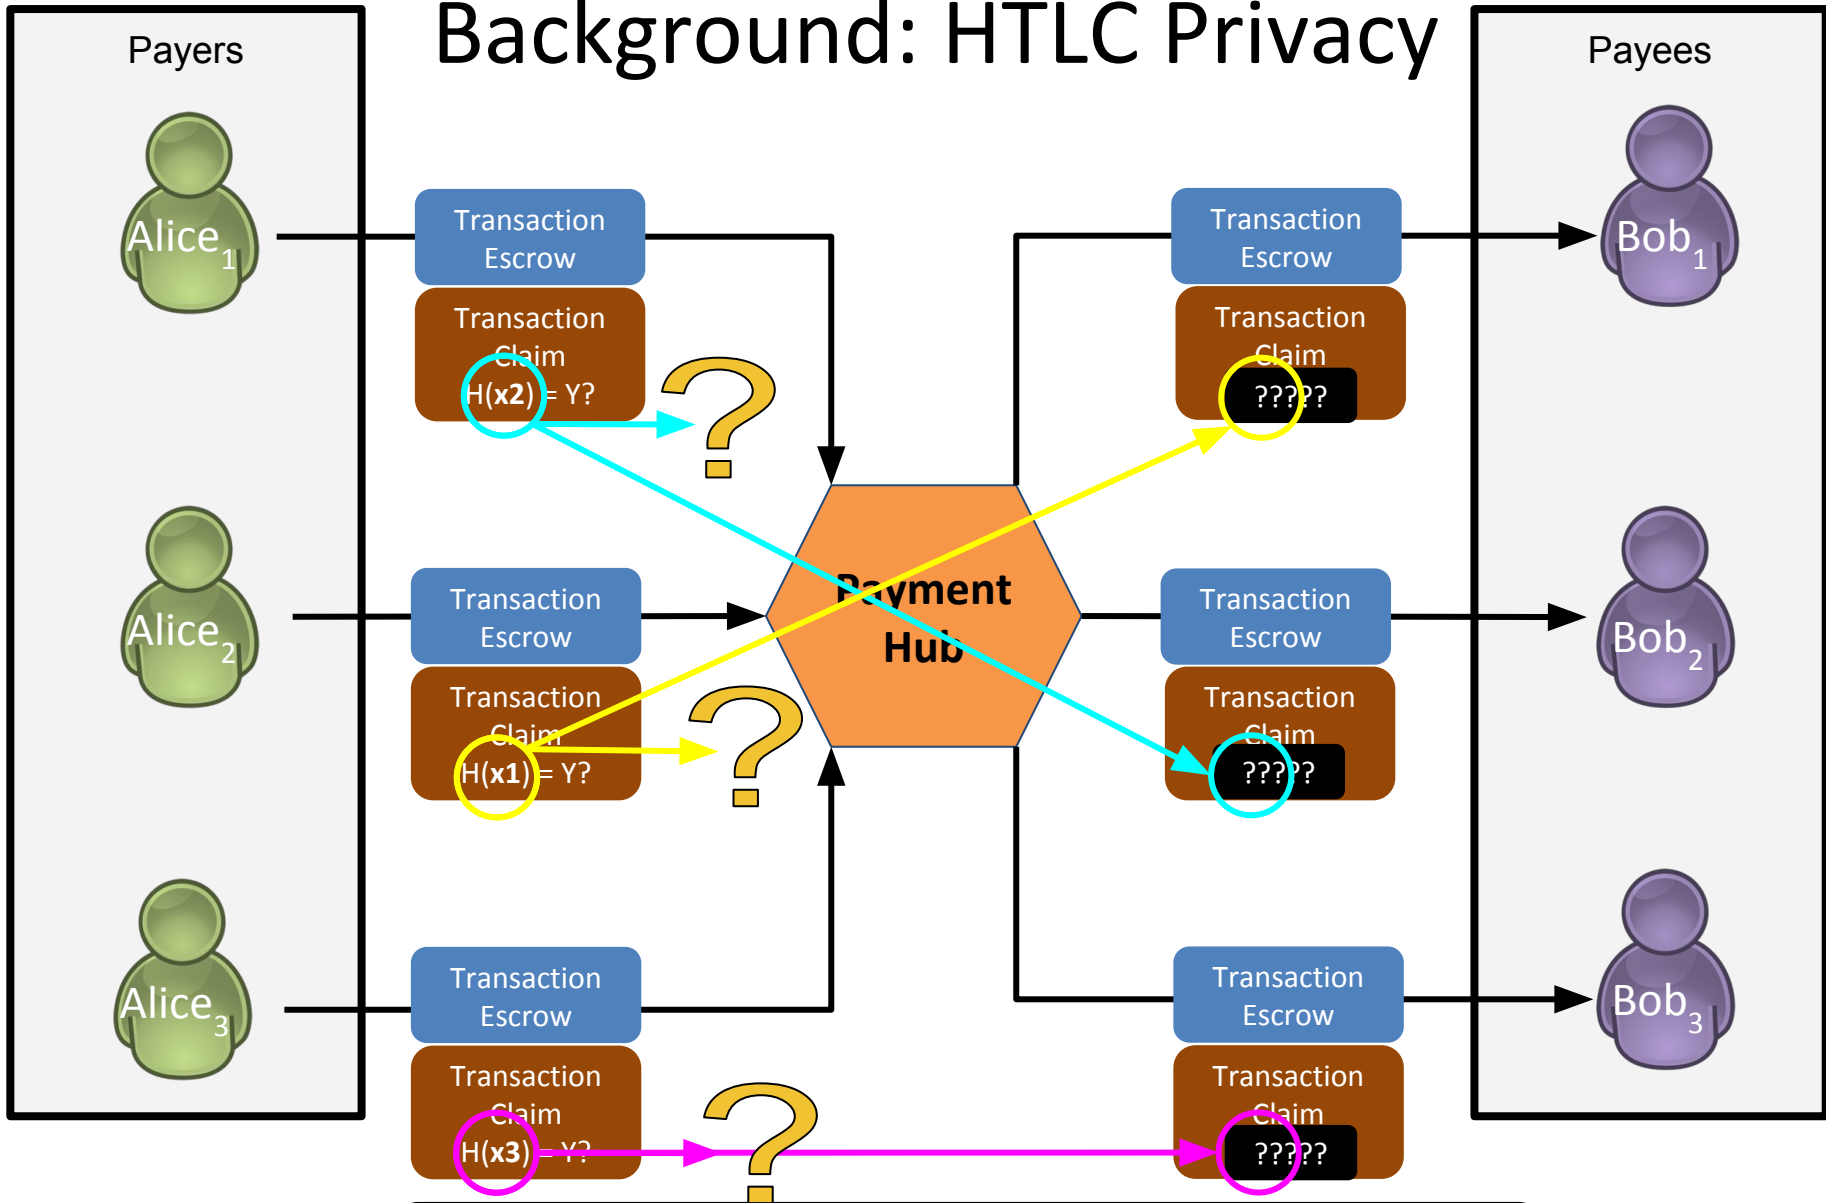
**Thus,** using hash locked transactions or HTLCs a payment hub can prevent theft, however this provides no privacy against the payment hub.

# Background: HTLC Privacy





# Background: HTLC Privacy

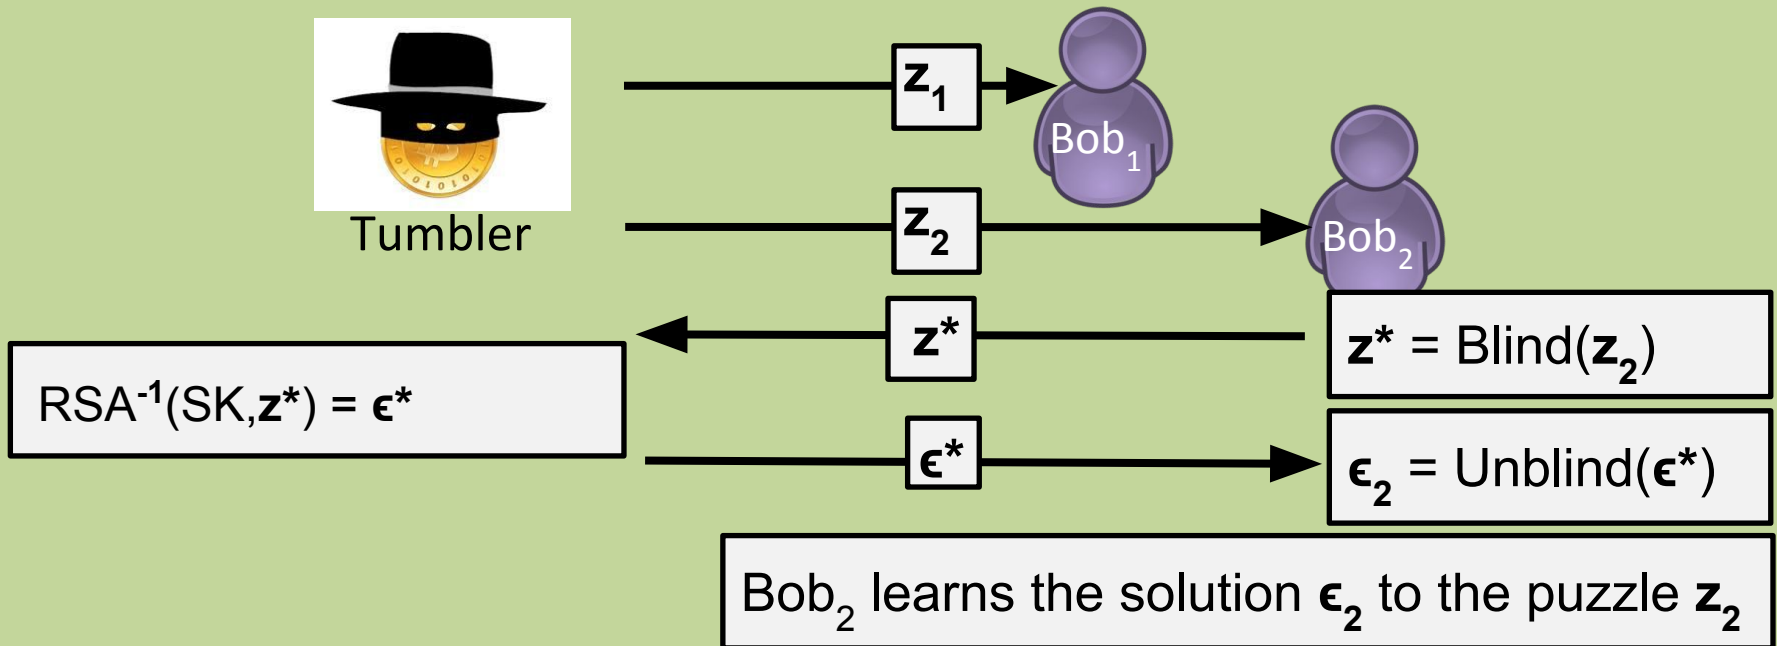


The main idea behind TumbleBit is a protocol which provides **atomicity** but is also **unlinkable** (i.e. private). Think of it like Unlinkable or Private HTLCs.

# Background: RSA Puzzles

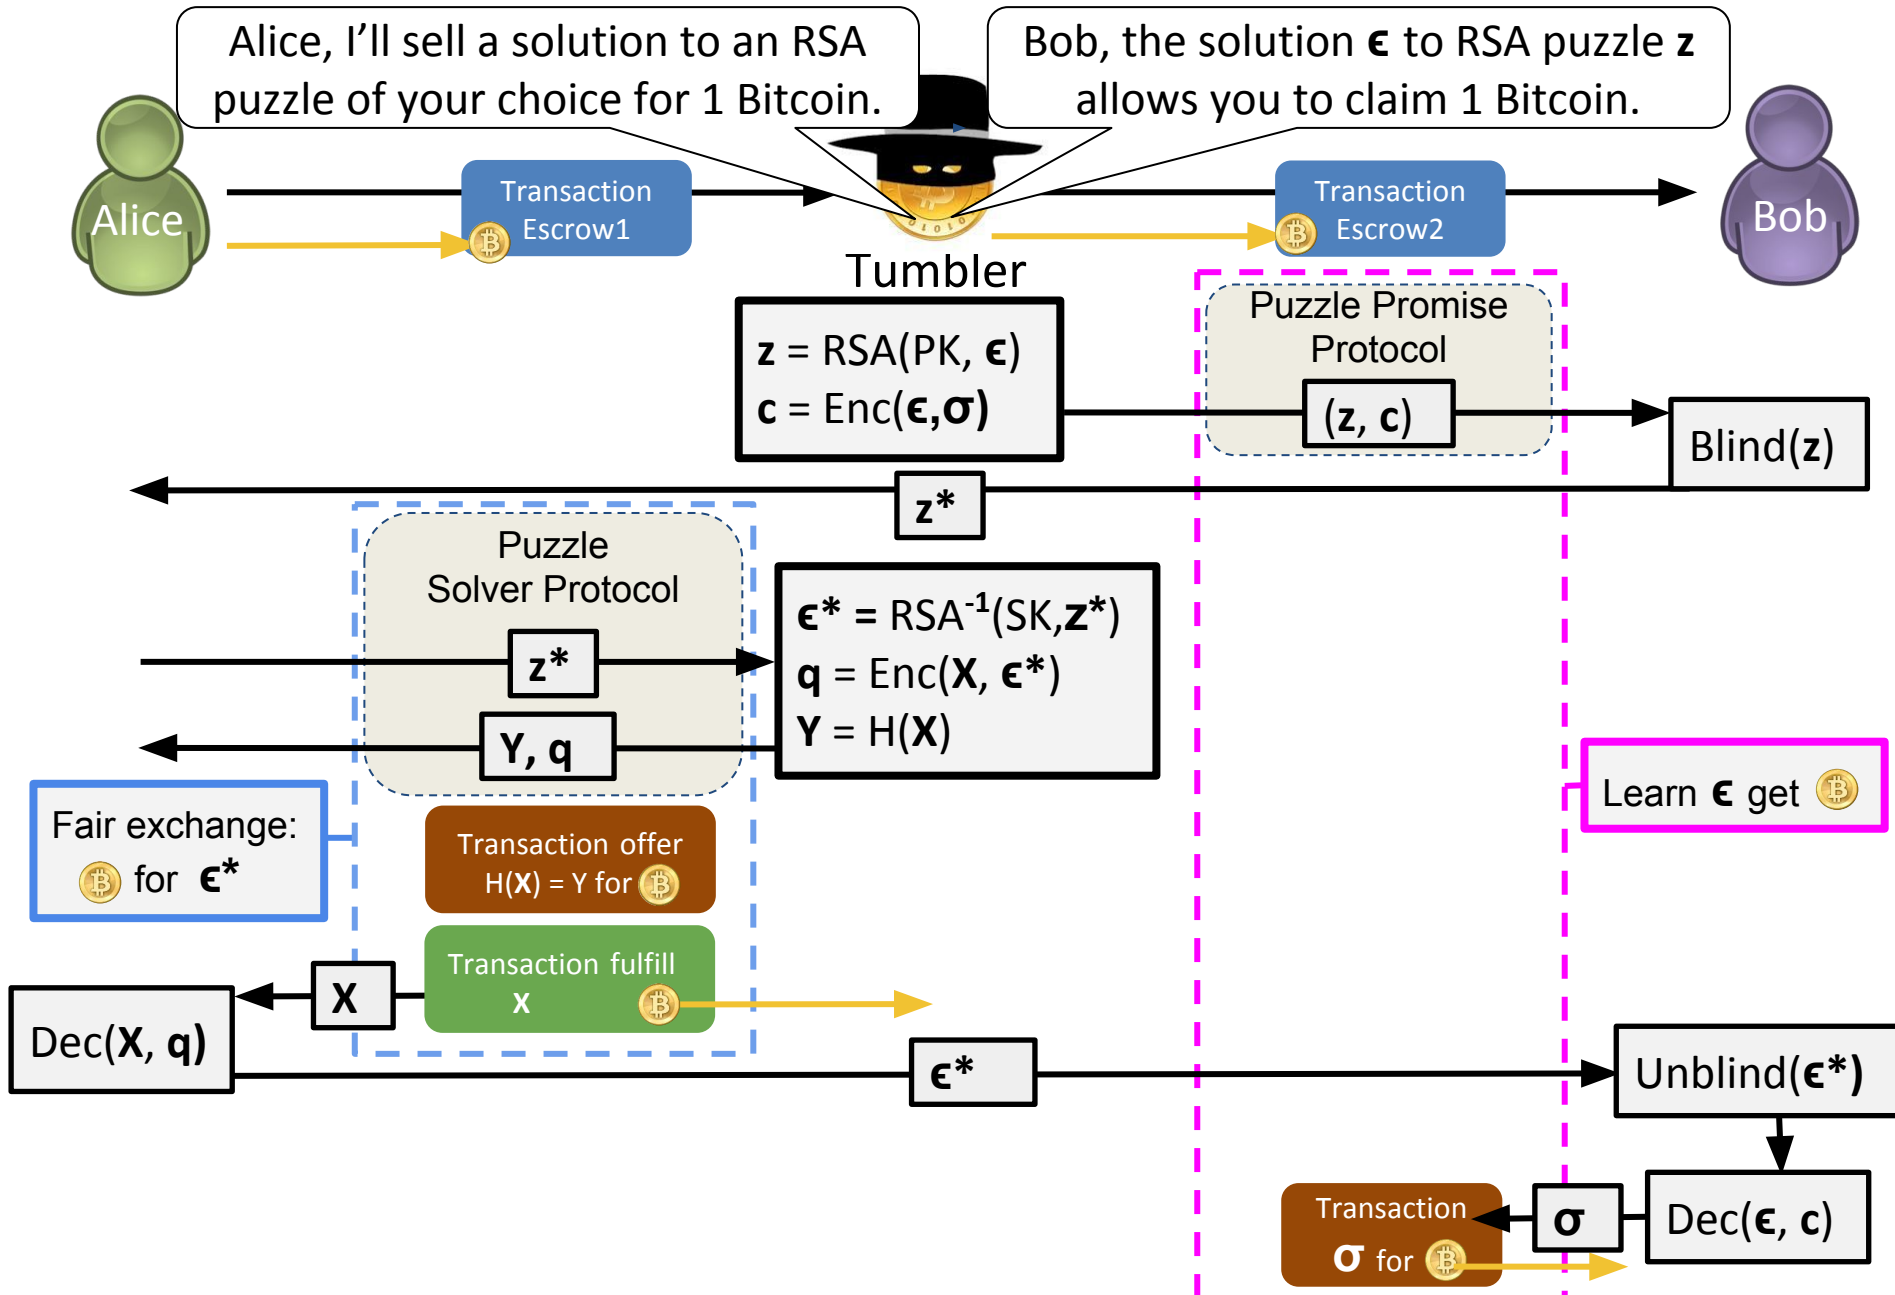
- An RSA Puzzle is just a “textbook RSA encryption” of some value  $\epsilon$ :  
$$\text{RSA}(\text{PK}, \epsilon) = z$$
- Only the party that knows SK can solve RSA puzzles:  
$$\text{RSA}^{-1}(\text{SK}, z) = \text{RSA}^{-1}(\text{SK}, \text{RSA}(\text{PK}, \epsilon)) = \epsilon$$

## RSA blinding can be used to blind RSA puzzles



Tumbler can not link the blinded RSA puzzle it solves  $z^*$  to any of the RSA puzzles it issued ( $z_1, z_2$ ).

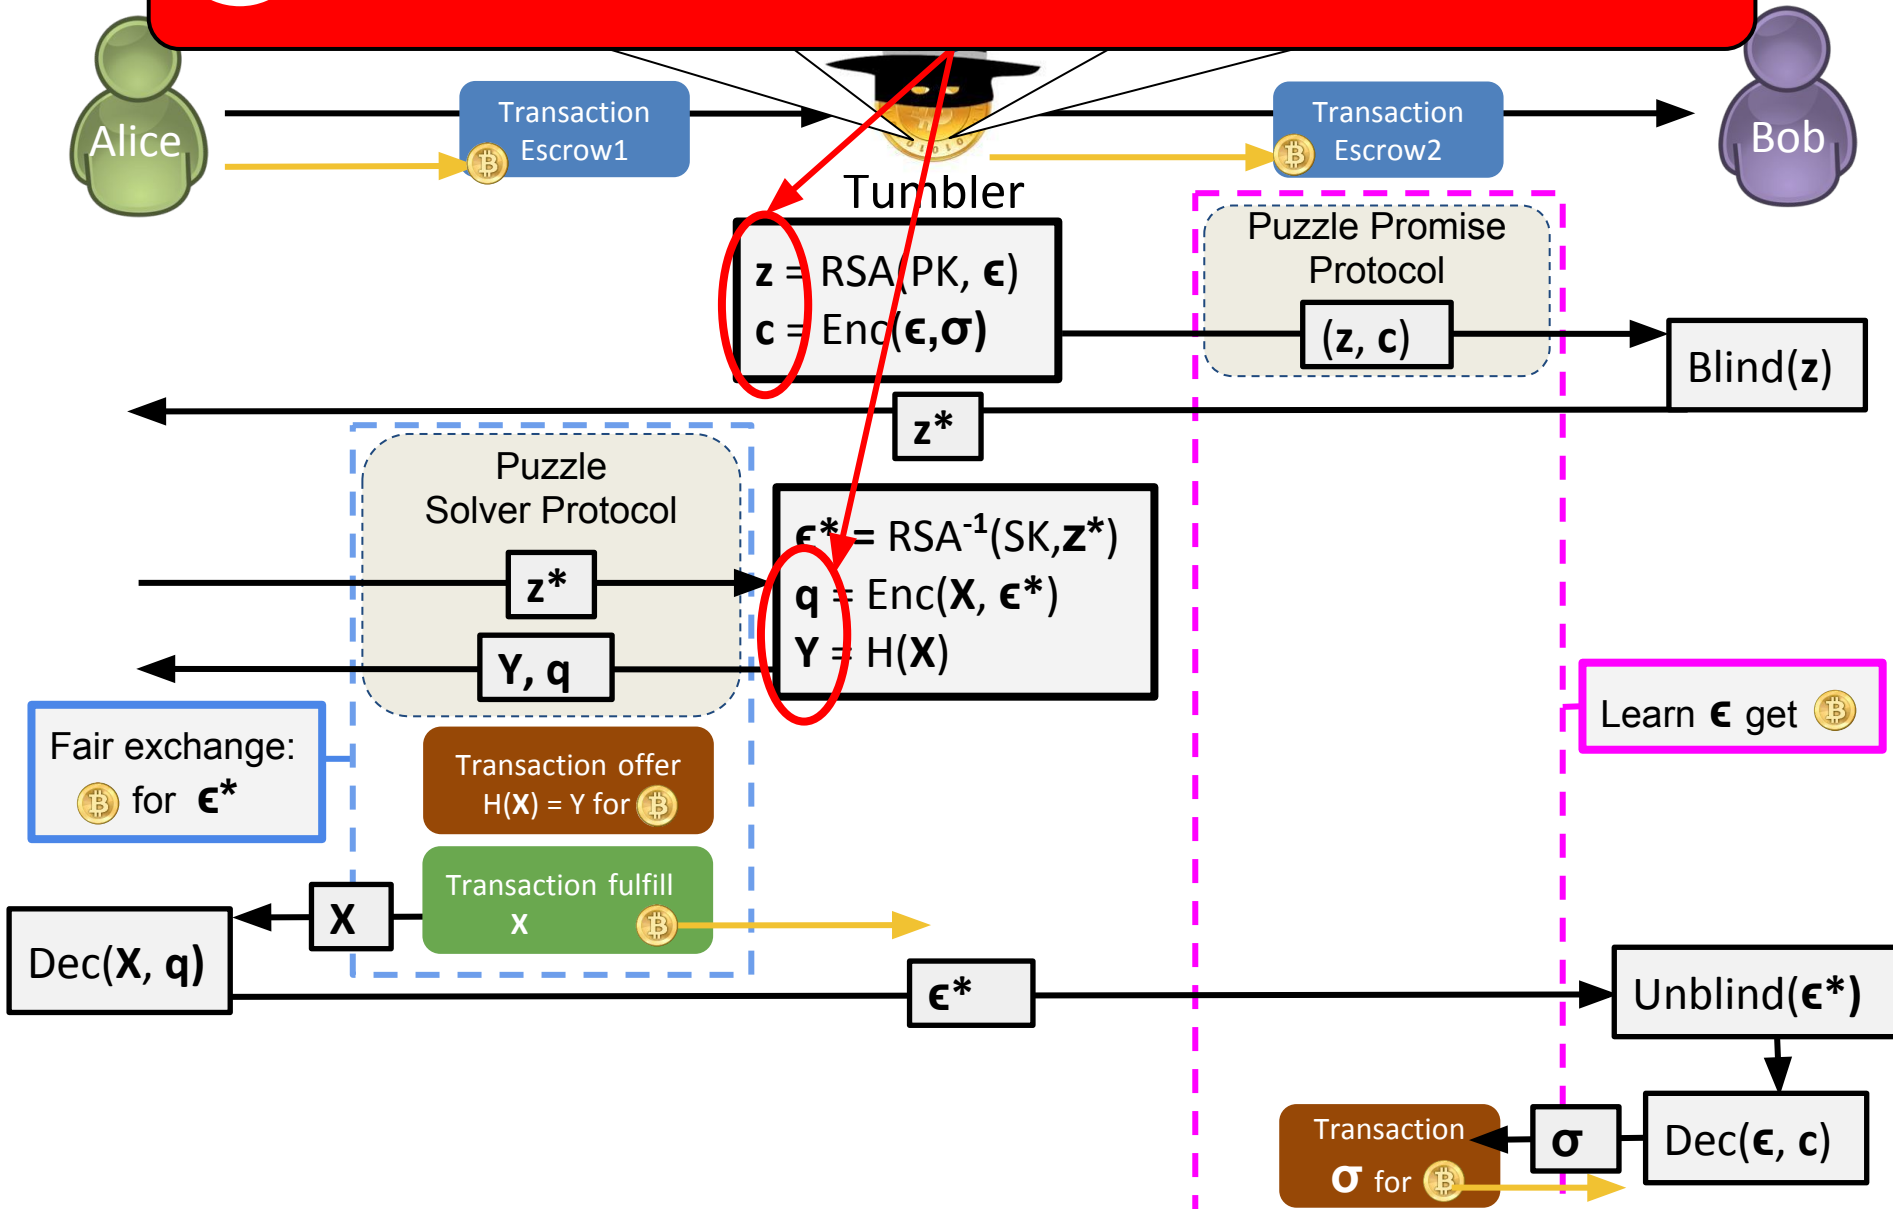
# TumbleBit: Protocol Overview



# TumbleBit: Protocol Overview



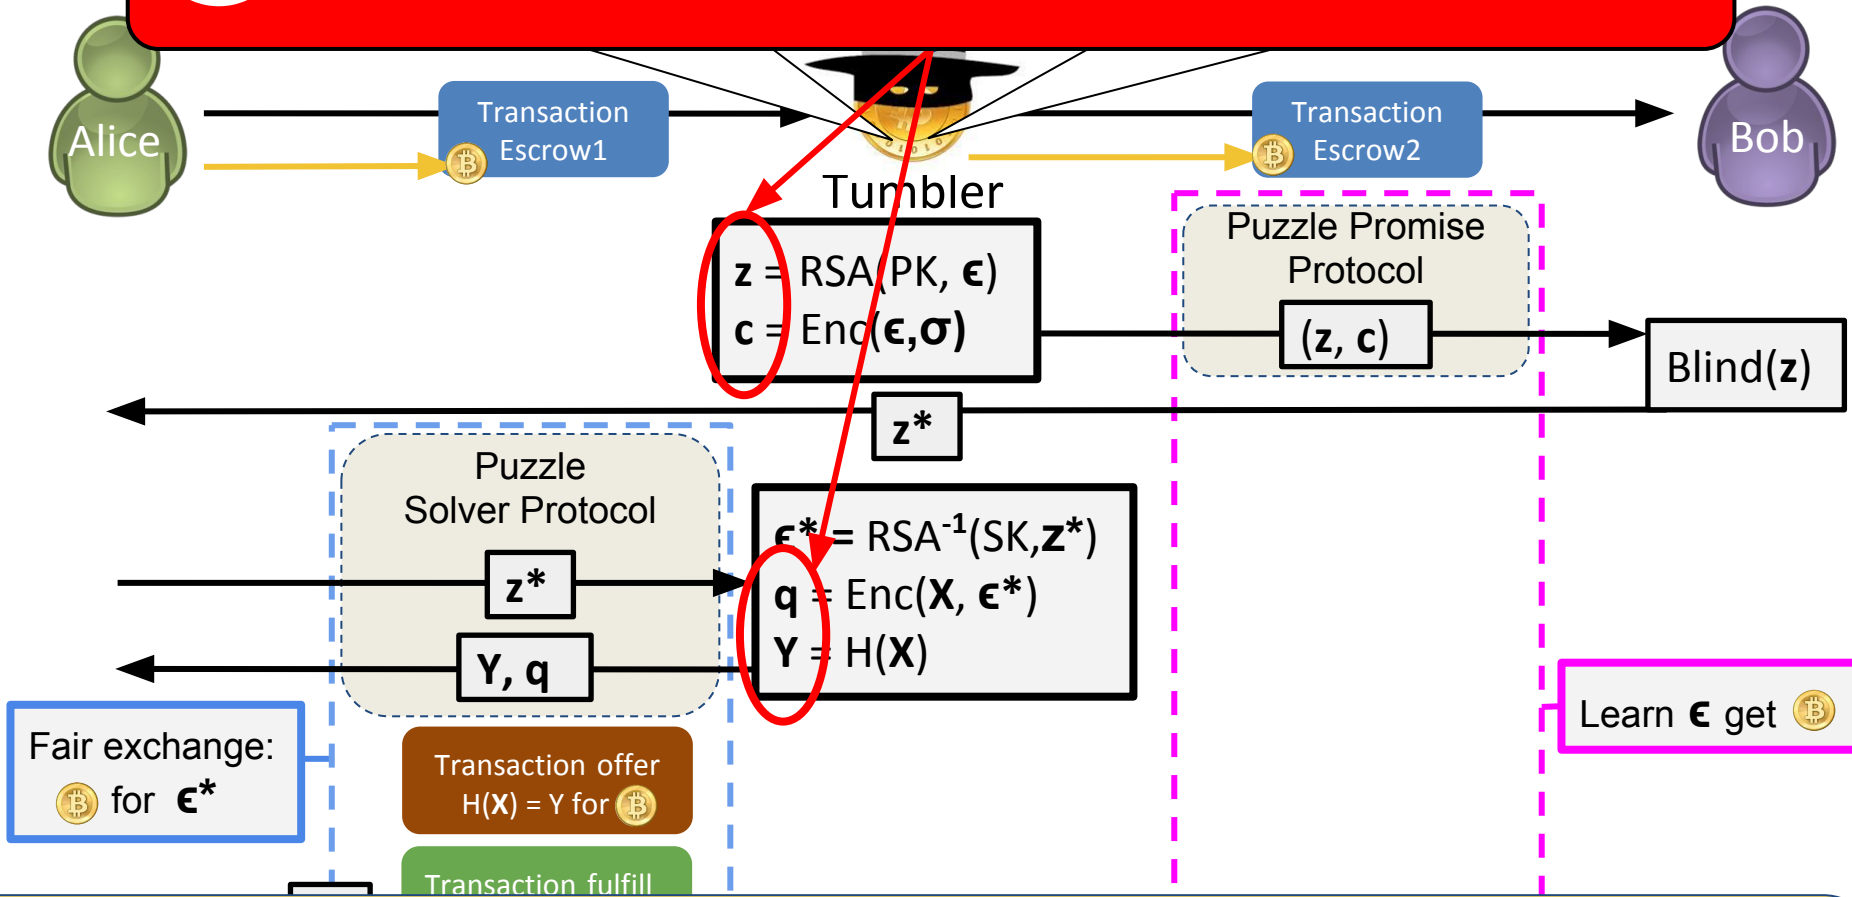
If Tumbler corrupts  $z$ ,  $c$ ,  $X$ , or  $q$  it can cheat Alice or Bob!



# TumbleBit: Protocol Overview



If Tumbler corrupts  $z$ ,  $c$ ,  $X$ , or  $q$  it can cheat Alice or Bob!



TumbleBit prevents this via two protocols:

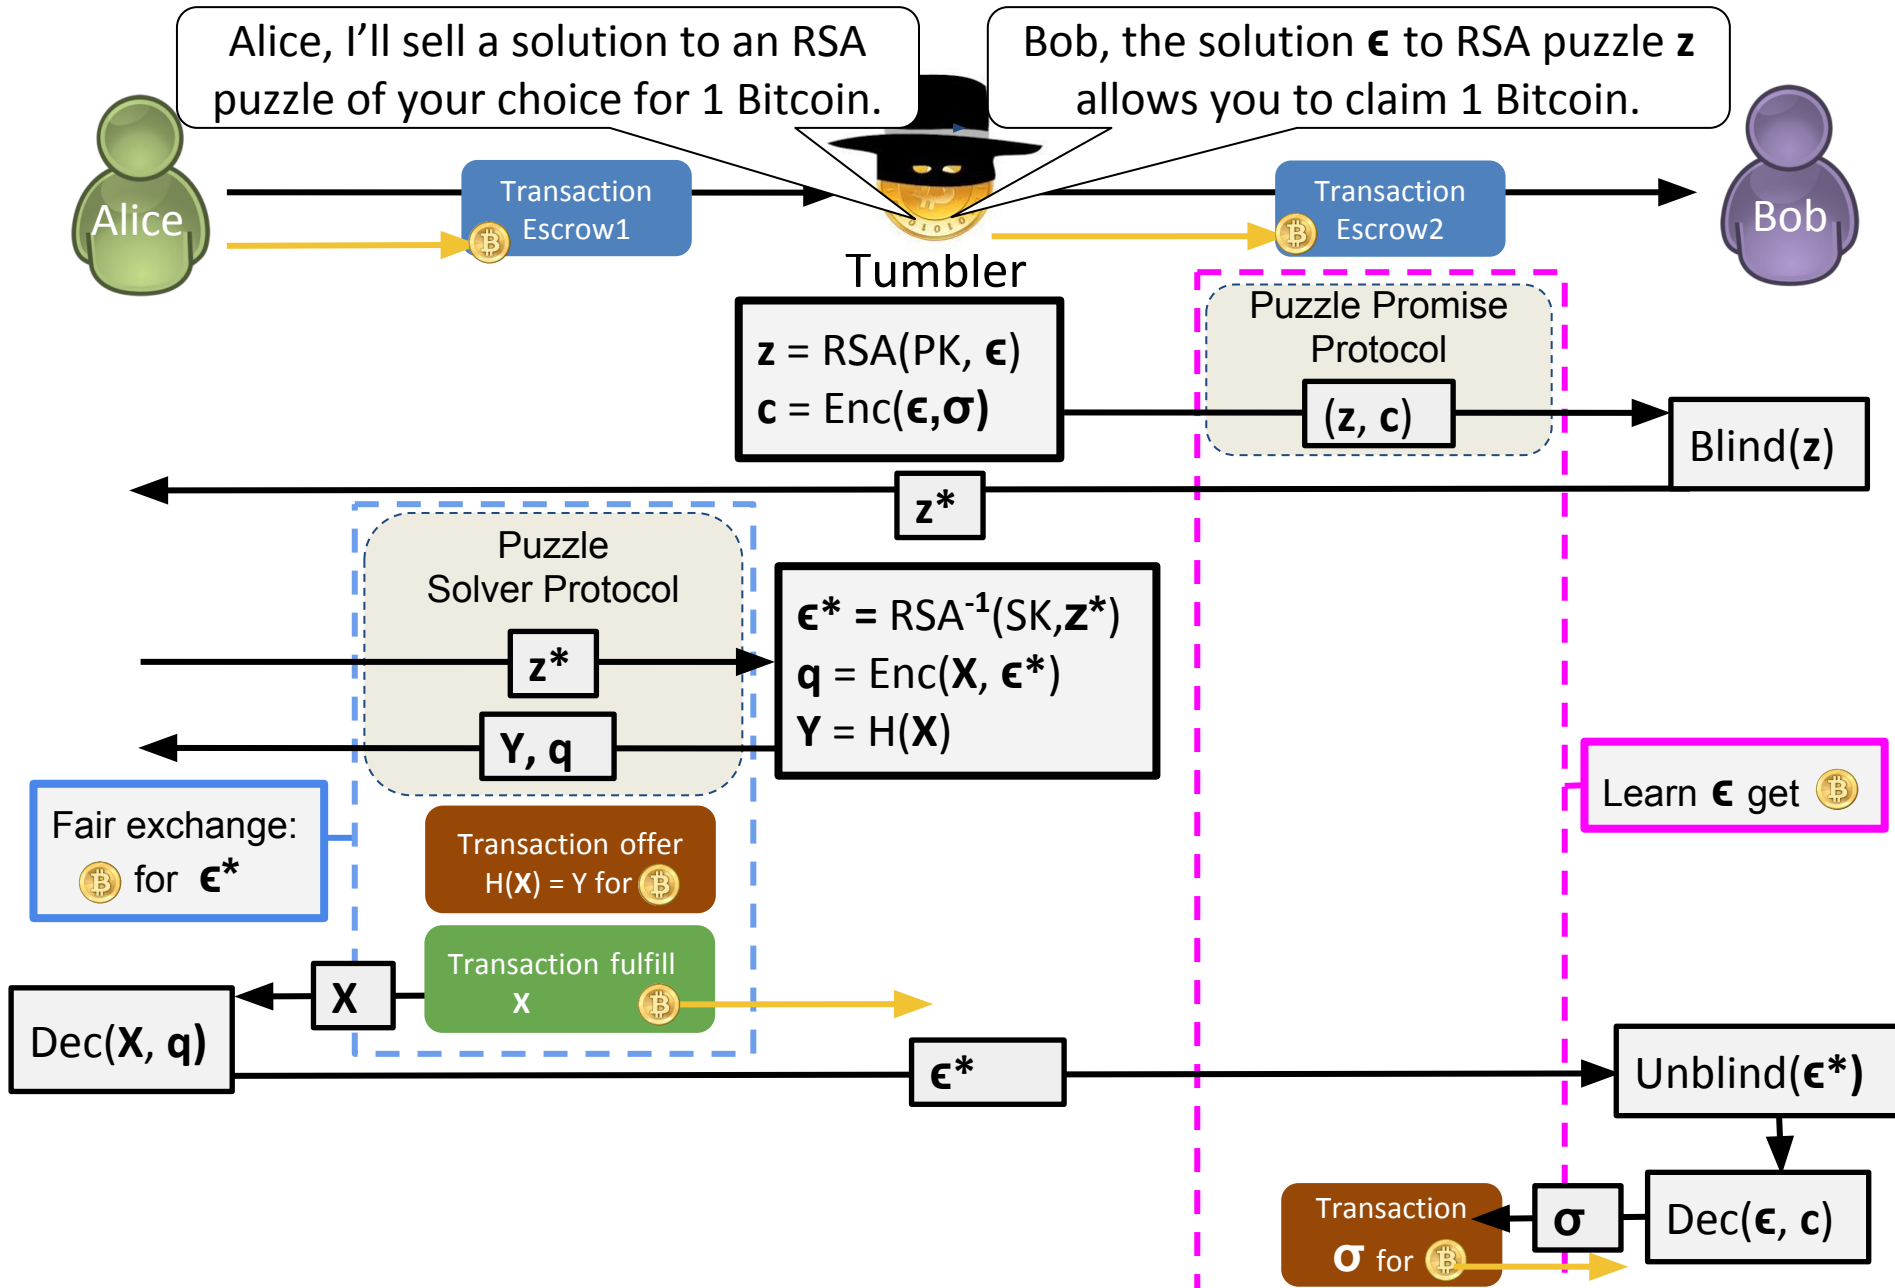
## Puzzle-Solver-Protocol:

Tumbler convinces Alice the preimage  $X$  where  $\text{Hash}(X) = Y$  will allow her to learn  $\epsilon^*$ .

## Puzzle-Promise-Protocol:

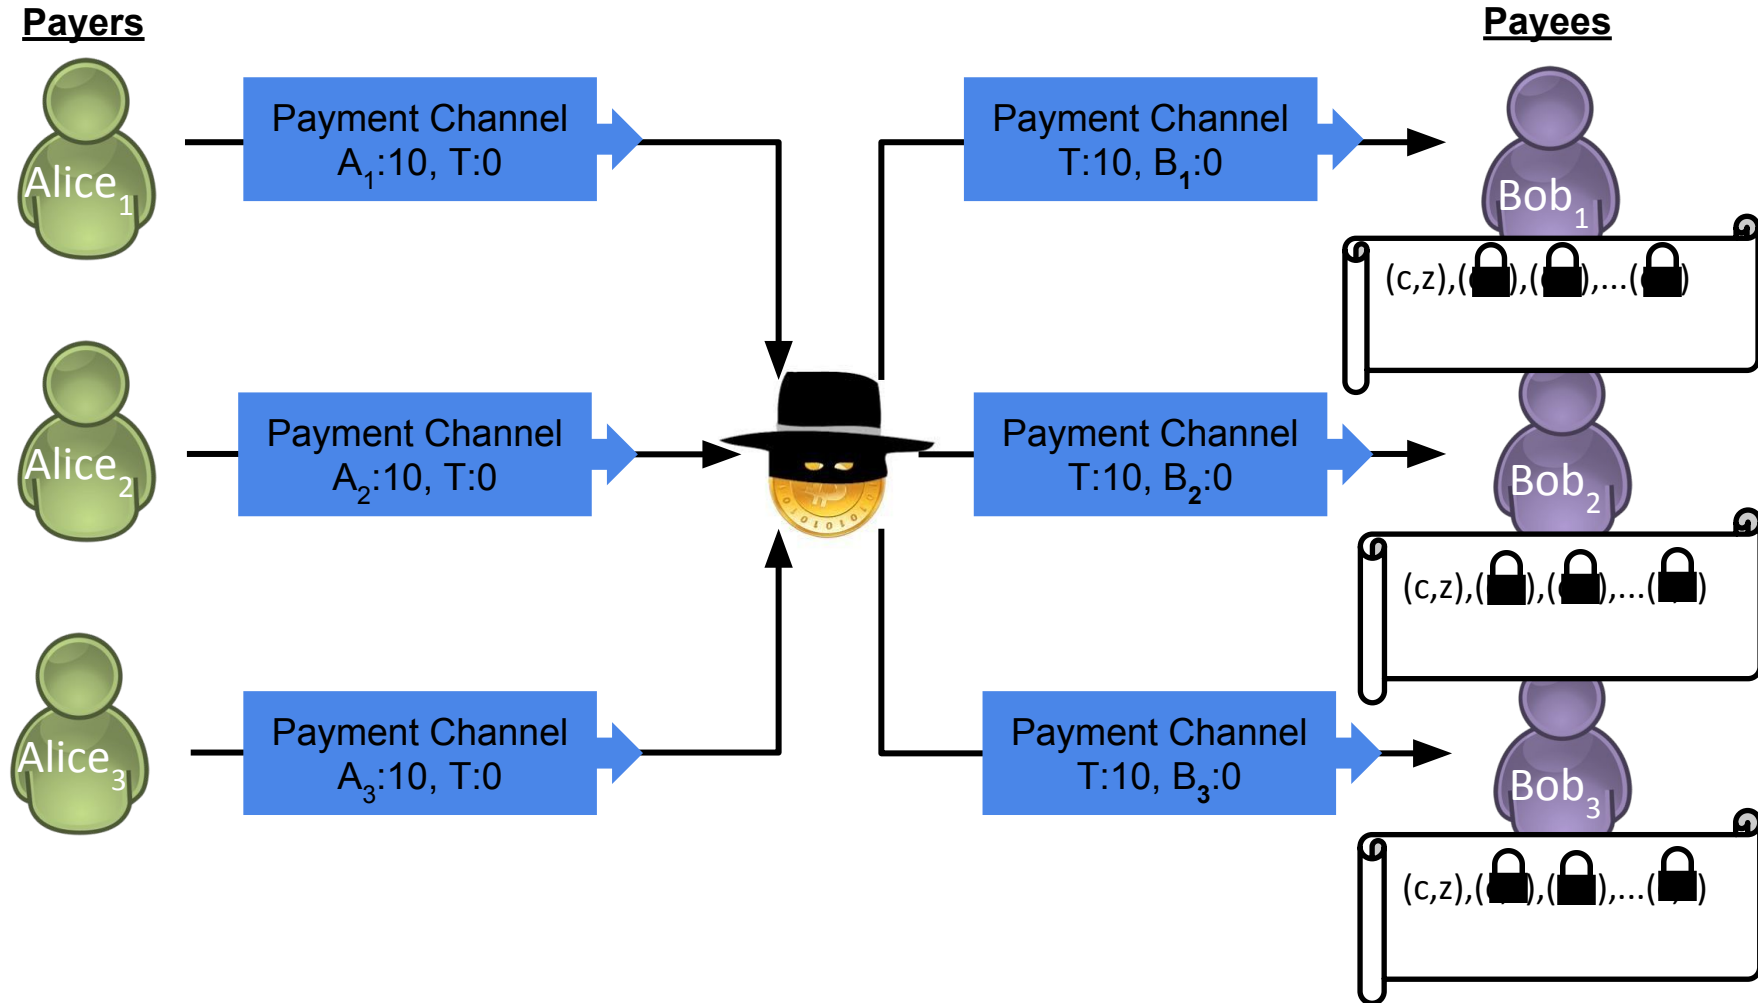
Tumbler convinces Bob that the solution to RSA puzzle  $z$  is a value  $\epsilon$  which allows him learn  $\sigma$ .

# TumbleBit: Protocol Overview



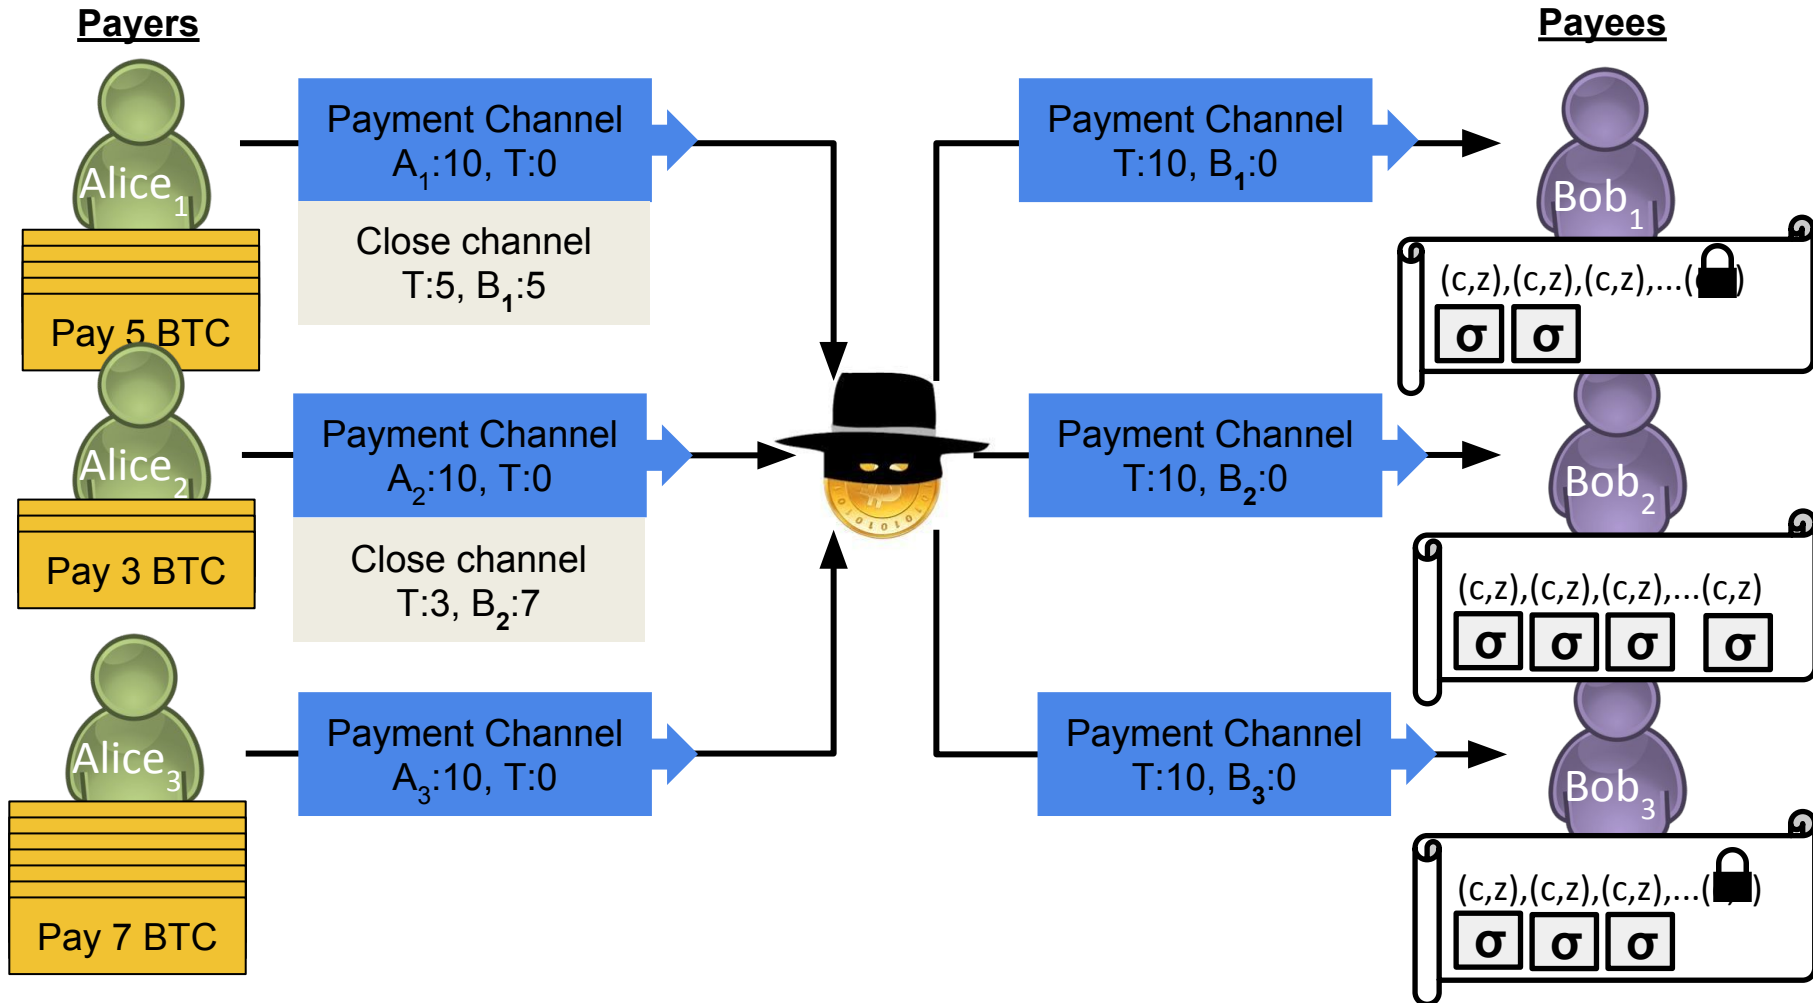
# TumbleBit: Phases

1. **Escrow Phase:** All payment channels setup.



# TumbleBit: Phases

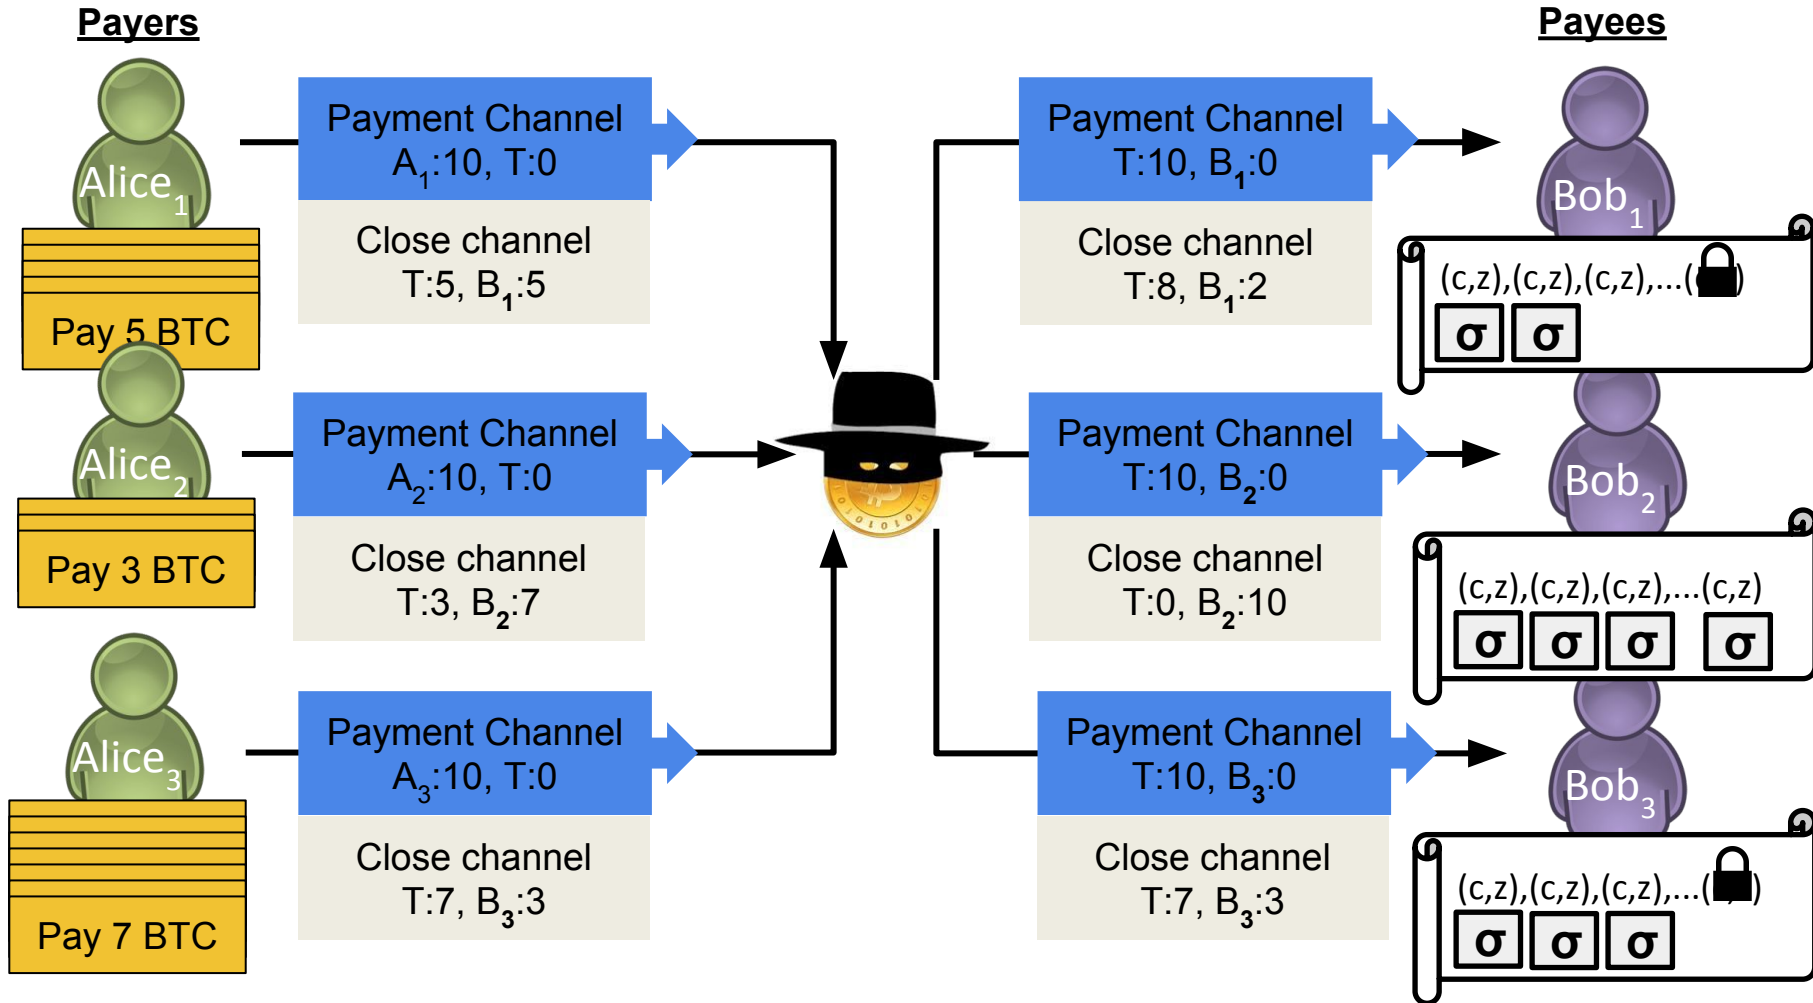
1. **Escrow Phase:** All payment channels setup.
2. **Payments Phase (~1 month):** Payers make payments.





# TumbleBit: Phases

1. **Escrow Phase:** All payment channels setup.
2. **Payments Phase (~1 month):** Payers make payments.
3. **Cashout Phase:** Payers and payees close their payment channels.



# TumbleBit: Phases


## Privacy offered the TumbleBit Payment Hub

### Tumbler's view:

(1) payer of each payment, (2) # of payments each payee received.

### Unlinkability def:


All interaction graphs compatible with the tumblers view are equally likely.



# TumbleBit: Classic Tumbler

**TumbleBit can also be a classic tumbler:**

Allows users to privately move bitcoins to an unlinked fresh address.

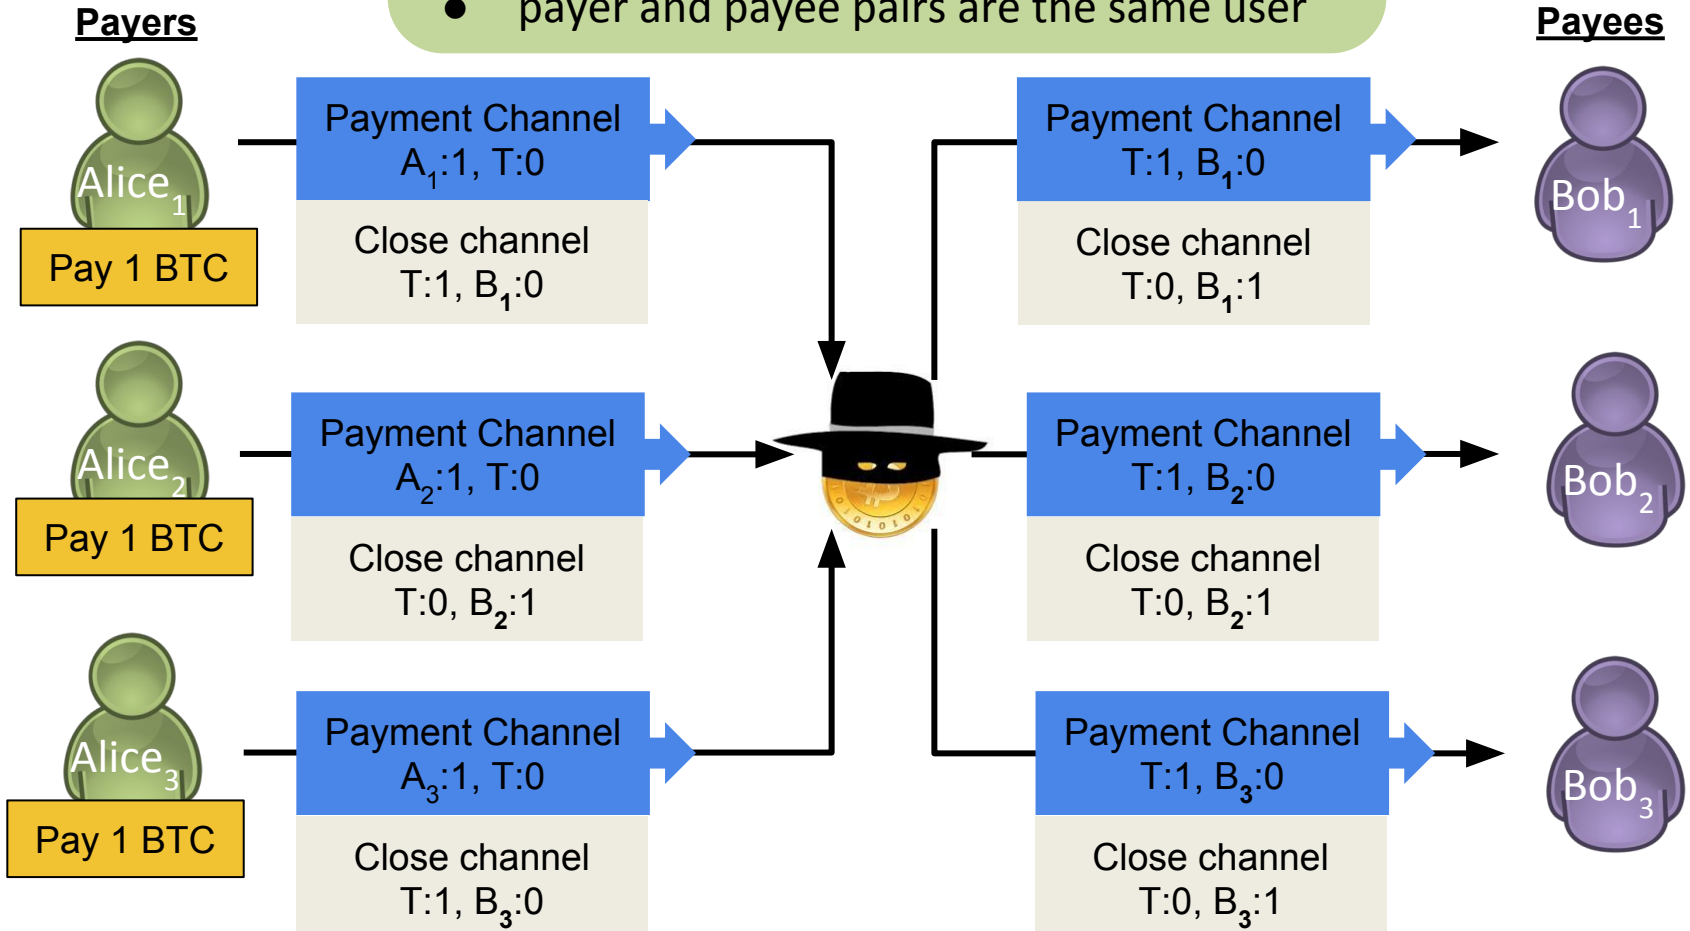


This is also sometimes known as a mixing service or mix.

# TumbleBit: Classic Tumbler

To run TumbleBit as a Classic Bitcoin Tumbler:

- Each payer just makes one payment.
- Each payee accepts only one payment.
- # of payers = # of payees.
- payer and payee pairs are the same user



**Provides k-anonymity:**

Where  $k = \# \text{ of payers} = \# \text{ of payee}$ .

# Compared to other Tumblers

**Vulnerable to DoS & Sybil Attacks**



**Limited Anonymity**



**TumbleBit**

**Mixing takes  
hours  
Xim**

**Vulnerable to bitcoin theft**



**Blindcoin:**



**Intermediary  
breaks  
anonymity**


# TumbleBit: Implementation

We wrote a proof-of-concept implementation of the Classic Tumbler:

- We are working on improving it and making it user friendly.
- Sourcecode and a development roadmap are available on

We “tumbled” 800 payments:

You can see the transactions on the mainnet blockchain.  
TXIDs available in our paper.



Our implementation is Performant (per TumbleBit payment):

- 326 KB of Bandwidth,
- Puzzle-Solver takes ~0.4 seconds to compute
- Total time depends on network latency:  
No latency ~0.6 seconds.  
Boston to Tokyo ~6 seconds (clear) and ~11 seconds  
...(both parties use TOR)

# Conclusion

**TumbleBit provides,  
private untrusted scalable payments via today's Bitcoin**

1. **Private:** Unlinkable or k-anonymous payments
2. **Trustless:** Tumbler can not steal or link payments.
3. **Scalable (payment hub):** scales Bitcoin's transaction velocity and volume.

**We have running code (for TumbleBit classic tumbler):**

- Our code runs on Bitcoin's mainnet blockchain.
- We have published our code on github.
- ...and we working to improve it and make TumbleBit easy and safe to use.

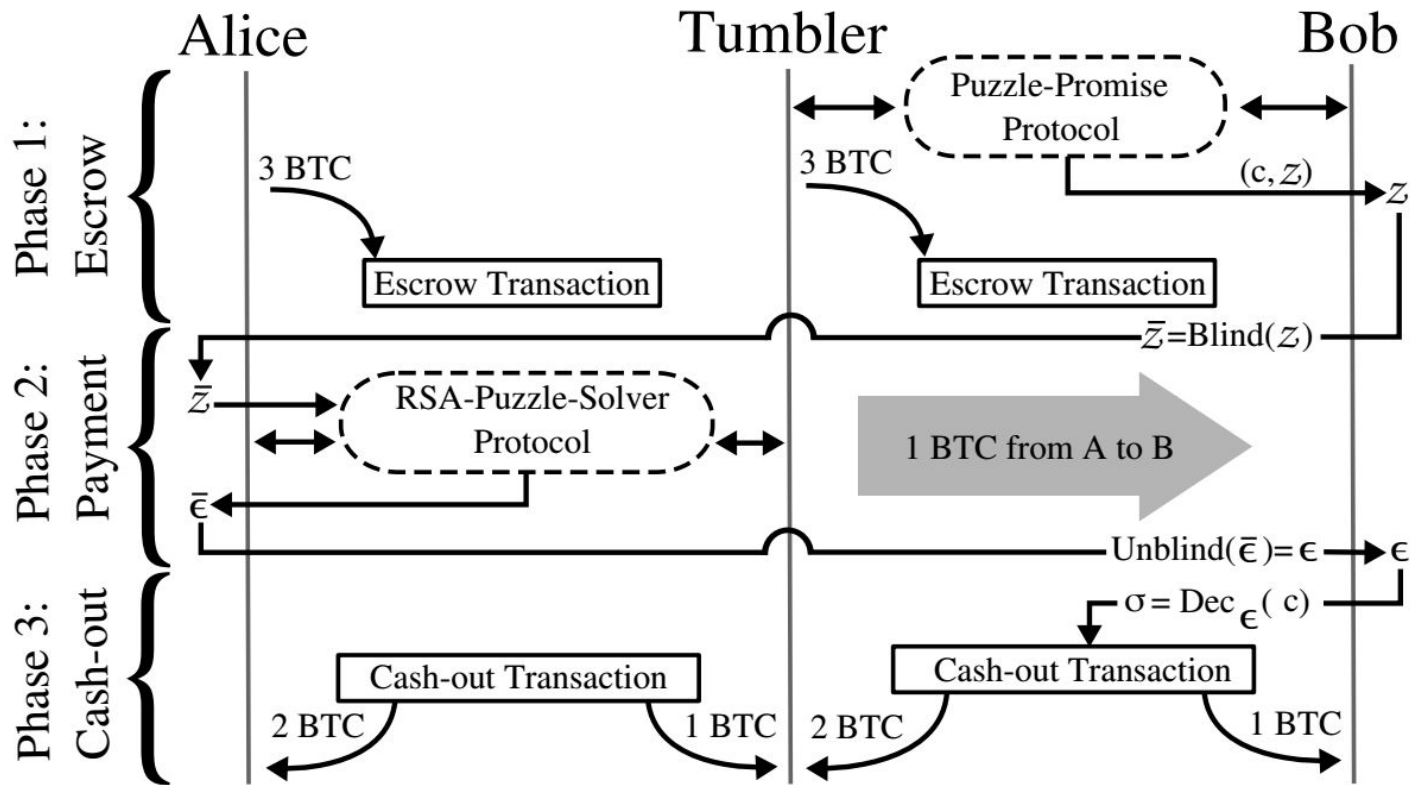


We are hiring a full time engineer (Boston),  
email me if interested.

# Questions?

Source code + roadmap: <https://github.com/BUSEC/TumbleBit>

Paper: <https://eprint.iacr.org/2016/575.pdf>



Ask questions on twitter: @Ethan\_Heilman



# TumbleBit: Puzzle-Solver-Prot

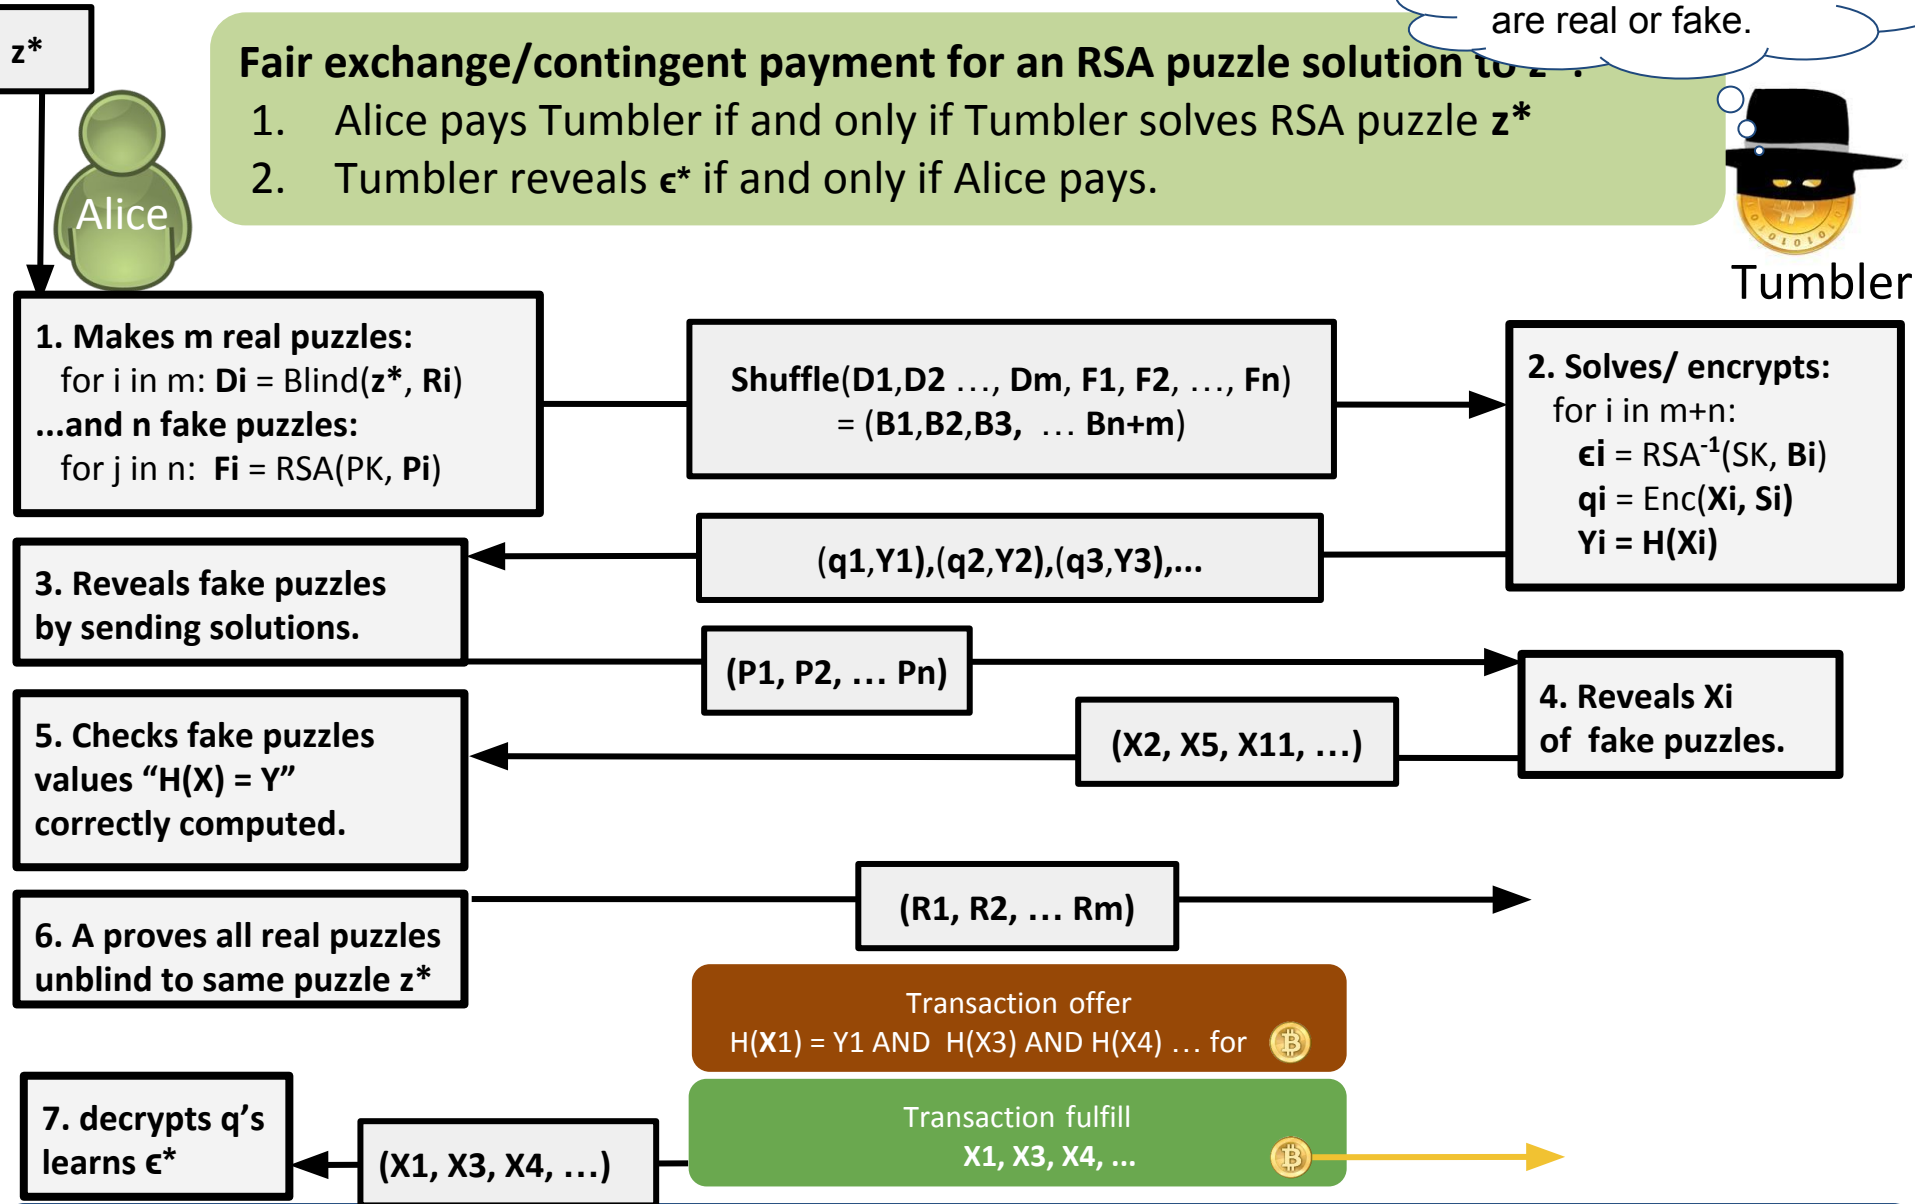
I can't tell which B's are real or fake.

Fair exchange/contingent payment for an RSA puzzle solution to  $z^*$ .

1. Alice pays Tumbler if and only if Tumbler solves RSA puzzle  $z^*$
2. Tumbler reveals  $\epsilon^*$  if and only if Alice pays.



Tumbler



If Tumbler computes any  $(q_i, \epsilon_i, Y_i)$  of the real puzzles correctly Alice learns  $\epsilon^*$ , thus to cheat Alice, Tumbler must corrupt all the real and none of the fake puzzles.

# TumbleBit: Puzzle-Solver-Prot

I can't tell which B's are real or fake.

Fair exchange/contingent payment for an RSA puzzle solution to  $z^*$ .

1. Alice pays Tumbler if and only if Tumbler solves RSA puzzle  $z^*$
2. Tumbler reveals  $\epsilon^*$  if and only if Alice pays.



$z^*$



1. Makes  $m$  real puzzles:  
for  $i$  in  $m$ :  $D_i = \text{Blind}(z^*, R_i)$   
...and  $n$  fake puzzles:  
for  $j$  in  $n$ :  $F_j = \text{RSA}(\text{PK}, P_j)$

Shuffle( $D_1, D_2, \dots, D_m, F_1, F_2, \dots, F_n$ )  
= ( $B_1, B_2, B_3, \dots, B_{m+n}$ )

2. Solves/ encrypts:  
for  $i$  in  $m+n$ :  
 $\epsilon_i = \text{RSA}^{-1}(\text{SK}, B_i)$   
 $q_i = \text{Enc}(X_i, S_i)$   
 $Y_i = H(X_i)$

( $a_1 Y_1$ ) ( $a_2 Y_2$ ) ( $a_3 Y_3$ ) ...

3. Reveals by sending

5. Checks values "correctly computed"

6. A proves all real puzzles unblind to same puzzle  $z^*$

7. decrypts  $q$ 's learns  $\epsilon^*$

Probability(Tumbler successfully cheats) =  $(m+n \text{ choose } m) = \sim 1/(2^{80})$   
 $m = \#$  of real puzzles = 15  
 $n = \#$  of fake puzzles = 285

( $R_1, R_2, \dots, R_m$ )

Transaction offer  
 $H(X_1) = Y_1$  AND  $H(X_3)$  AND  $H(X_4)$  ... for

Transaction fulfill  
 $X_1, X_3, X_4, \dots$

( $X_1, X_3, X_4, \dots$ )

If Tumbler computes any ( $q_i, \epsilon_i, Y_i$ ) of the real puzzles correctly Alice learns  $\epsilon^*$ , thus to cheat Alice, Tumbler must corrupt all the real and none of the fake puzzles.

# TumbleBit: Puzzle-Promise-Protocol

**At the end of this protocol:** Bob should be convinced that for a  $(z, c)$ :

1. The ciphertext  $c$  decrypts to  $\sigma$  under a key  $\epsilon$  i.e  $\text{Dec}(\epsilon, c) = \sigma$
2. **AND** the key  $\epsilon$  is the solution to the RSA-puzzle  $z$ .

**The protocol should never:** allow Bob to learn a valid  $\sigma$  (without paying).



This is why the protocol is hard,  
otherwise Tumbler could convince Bob  
by just sending  $(c, z, \epsilon, \sigma)$  and let Bob check.

# TumbleBit: Puzzle-Promise-Protocol

At the end of this protocol: Bob should be convinced that for a  $(z, c)$ :

1. The ciphertext  $c$  decrypts to  $\sigma$  under a key  $\epsilon$  i.e  $\text{Dec}(\epsilon, c) = \sigma$
2. **AND** the key  $\epsilon$  is the solution to the RSA-puzzle  $z$ .

The protocol should **never**: allow Bob to learn a valid  $\sigma$  (without paying).



1. B sends: a mix of hashes of valid and invalid claim transactions.

$B = H(T1), H(\text{invalid}), H(\text{invalid}), H(T4), H(\text{invalid}), H(T6)$

2. T Signs & Encrypts  $\sigma$ :

for  $B_i$  in  $B$ :

$$\sigma_i = \text{Sign}(B_i)$$

$$z_i = \text{RSA}^{-1}(\text{SK}, \epsilon_i), c_i = \text{Enc}(\epsilon_i, \sigma_i)$$

This is why the protocol is hard,  
otherwise Tumbler could convince Bob  
by just sending  $(c, z, \epsilon, \sigma)$  and let Bob check.

$T1, \text{invalid}, \text{invalid}, T4, \text{invalid}, T6$

3. B: reveals transactions.

4. T Reveals:  $\epsilon_i$  for invalid transactions.

$\epsilon_2, \epsilon_3, \epsilon_5$

5. B checks: invalid transactions  $\sigma_i$  are correctly computed.

6. Bob and Tumbler run "quotient protocol" ensuring that:  
if Bob learns  $\epsilon_1$ , Bob can use that knowledge to learn  $\epsilon_4, \epsilon_6$ .  
 $(\epsilon_4/\epsilon_1 \bmod N, \epsilon_6/\epsilon_4 \bmod N)$

If Tumbler computes any  $(\epsilon_i, \sigma_i)$  of the valid transactions correctly Bob learns a  $\sigma$ /gets paid, **thus** to cheat Bob, Tumbler must all corrupt all the valid and none of the invalid transactions.

# TumbleBit: Puzzle-Promise-Protocol

At the end of this protocol: Bob should be convinced that for a  $(z, c)$ :

1. The ciphertext  $c$  decrypts to  $\sigma$  under a key  $\epsilon$  i.e  $\text{Dec}(\epsilon, c) = \sigma$
2. AND the key  $\epsilon$  is the solution to the RSA-puzzle  $z$ .

The protocol should never: allow Bob to learn a valid  $\sigma$  (without paying).



1. B sends: a mix of hashes of valid and invalid claim transactions.

$B = H(T1), H(\text{invalid}), H(\text{invalid}), H(T4), H(\text{invalid}), H(T6)$

2. T Signs & Encrypts  $\sigma$ :  
for  $B_i$  in  $B$ :  
 $\sigma_i = \text{Sig}(\dots)$   
 $z_i = \text{RS}(\dots)$

This is why the protocol is hard,  
otherwise Tumbler could convince Bob

Probability(Tumbler successfully cheats) =  $(m+n \text{ choose } m) = \sim 1/(2^{80})$   
 $m = \# \text{ of valid transactions} = 42$   
 $n = \# \text{ of invalid transactions} = 42$

4. T Reveals:  $\epsilon_i$  for invalid transactions.

$\epsilon_2, \epsilon_3, \epsilon_5$

3. B checks: invalid transactions  $\sigma_i$  are correctly computed.

6. Bob and Tumbler run "quotient protocol" ensuring that:  
if Bob learns  $\epsilon_1$ , Bob can use that knowledge to learn  $\epsilon_4, \epsilon_6$ .  
 $(\epsilon_4/\epsilon_1 \text{ mod } N, \epsilon_6/\epsilon_4 \text{ mod } N)$

If Tumbler computes any  $(\epsilon_i, \sigma_i)$  of the valid transactions correctly Bob learns a  $\sigma$ /gets paid, thus to cheat Bob, Tumbler must all corrupt all the valid and none of the invalid transactions.