

Improved Straight-Line Extraction in the Random Oracle Model with Applications to Signature Aggregation

Yashvanth Kondi

abhi shelat



This Work

- We explore two dimensions of Fischlin's NIZKPoK compiler:
 - **Applicability:**
Only proven for Sigma protocols with 'quasi-unique responses'
(doesn't include logical OR, Pedersen commitment PoK, etc.)
Folklore: "works anyway"
 - 1a) Contrary to folklore: attack on Witness Indistinguishability
 - 1b) Simple randomization fixes the problem
 - **Computation cost:**
Usually the bottleneck — can we improve on it?
 - 2) Lower bound: Fischlin05 is optimal up to a small constant
 - 3) Application-specific optimization: 200× for EdDSA aggregation

Recap: Σ Protocol for Relation R

[Damgård 02]



$P(X, w)$



$V(X)$

Recap: Σ Protocol for Relation R

[Damgård 02]



$P(X, w)$



$V(X)$

Commitment

a

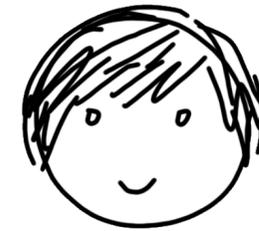


Recap: Σ Protocol for Relation R

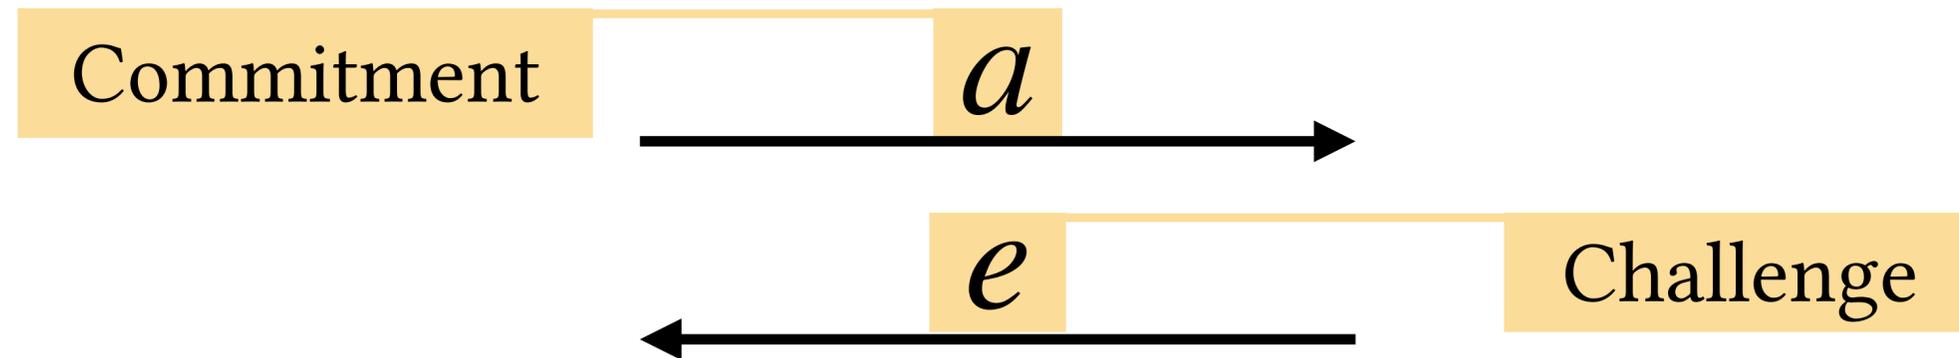
[Damgård 02]



$P(X, w)$

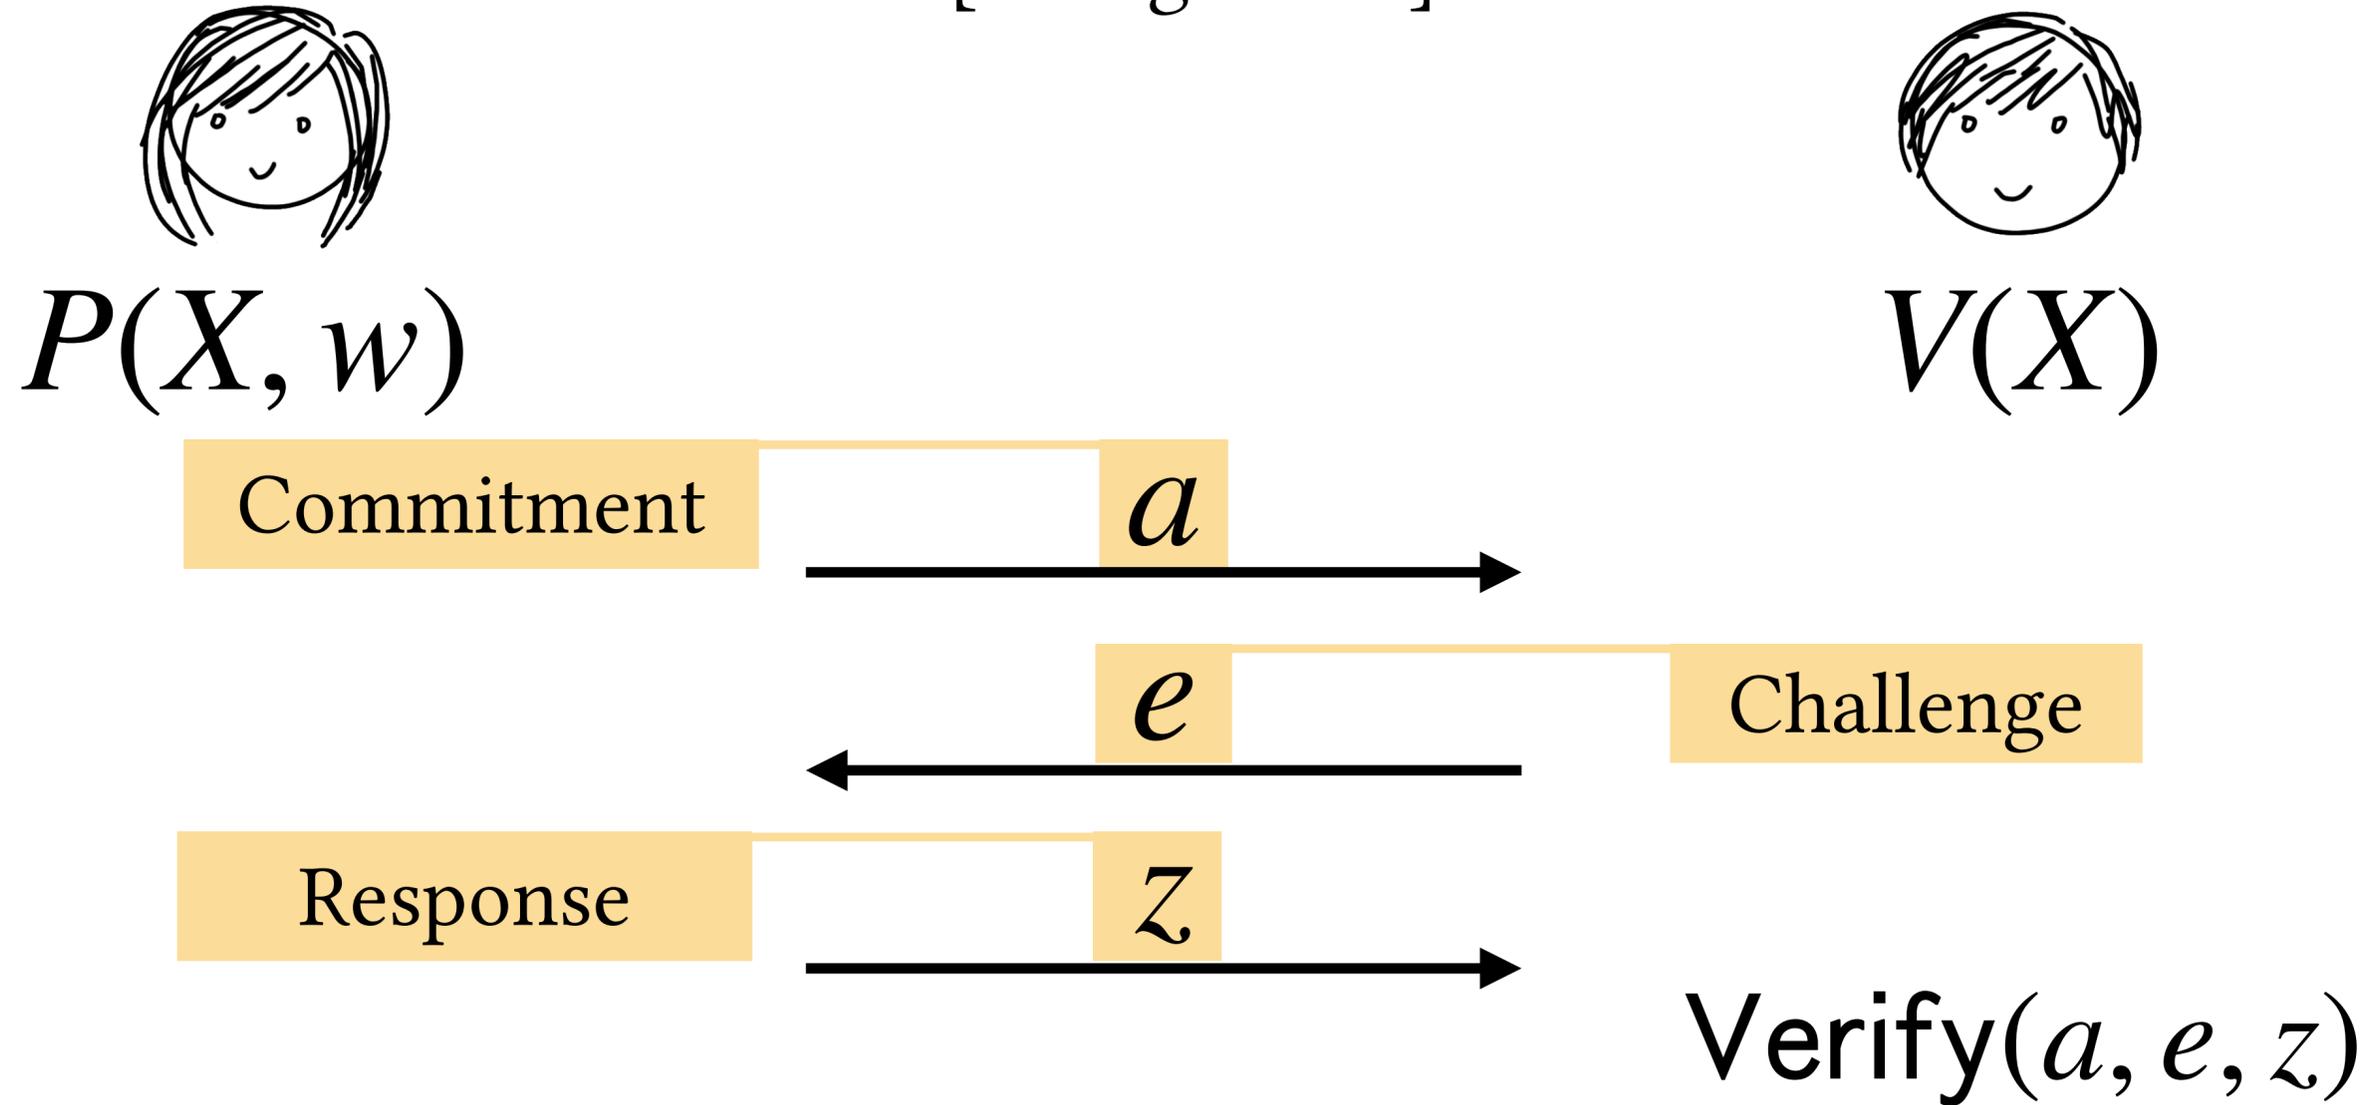


$V(X)$



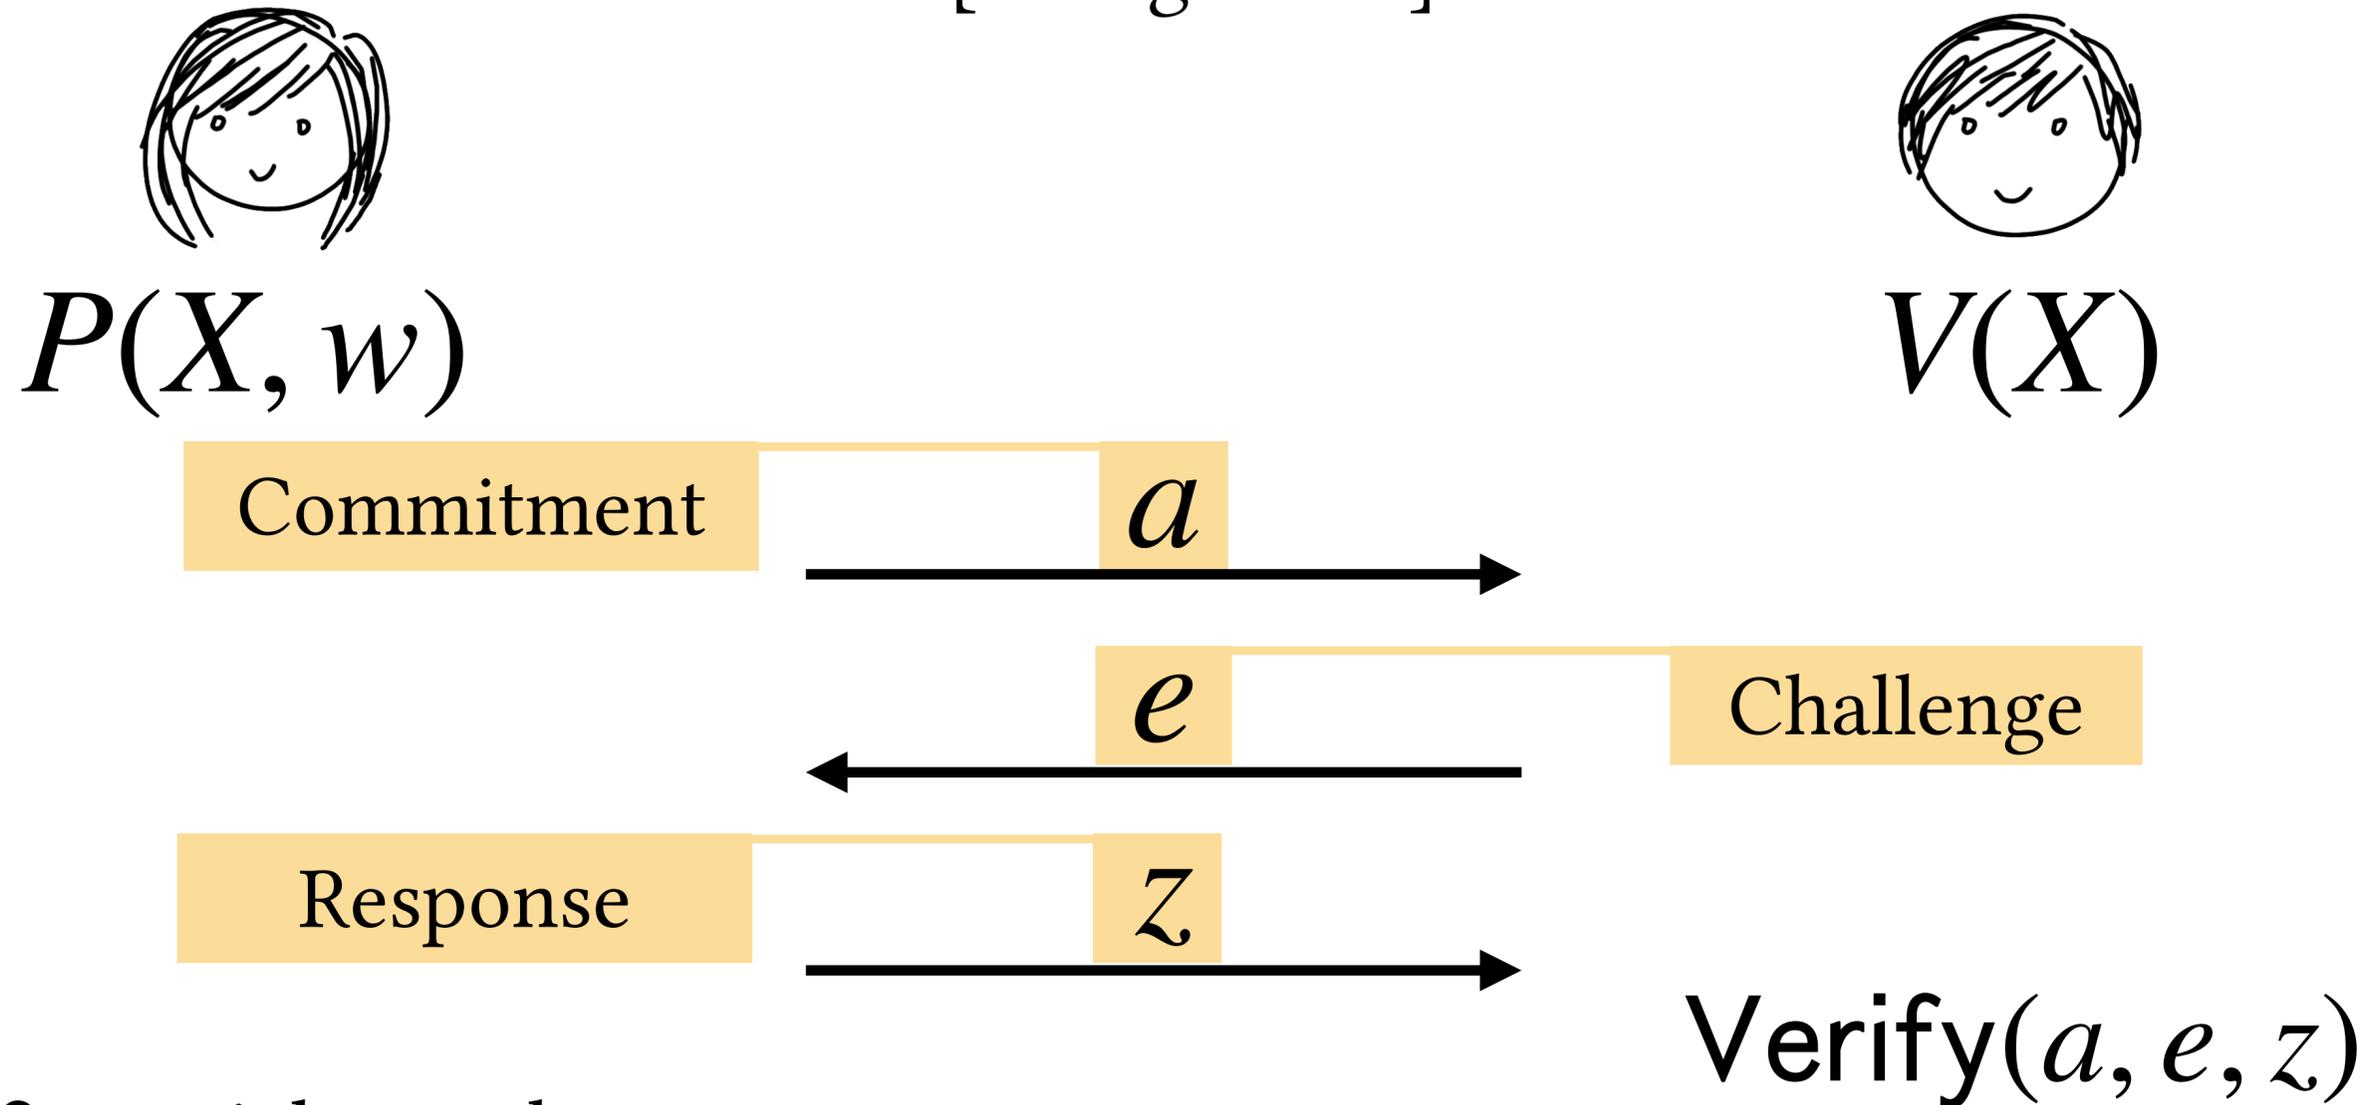
Recap: Σ Protocol for Relation R

[Damgård 02]



Recap: Σ Protocol for Relation R

[Damgård 02]

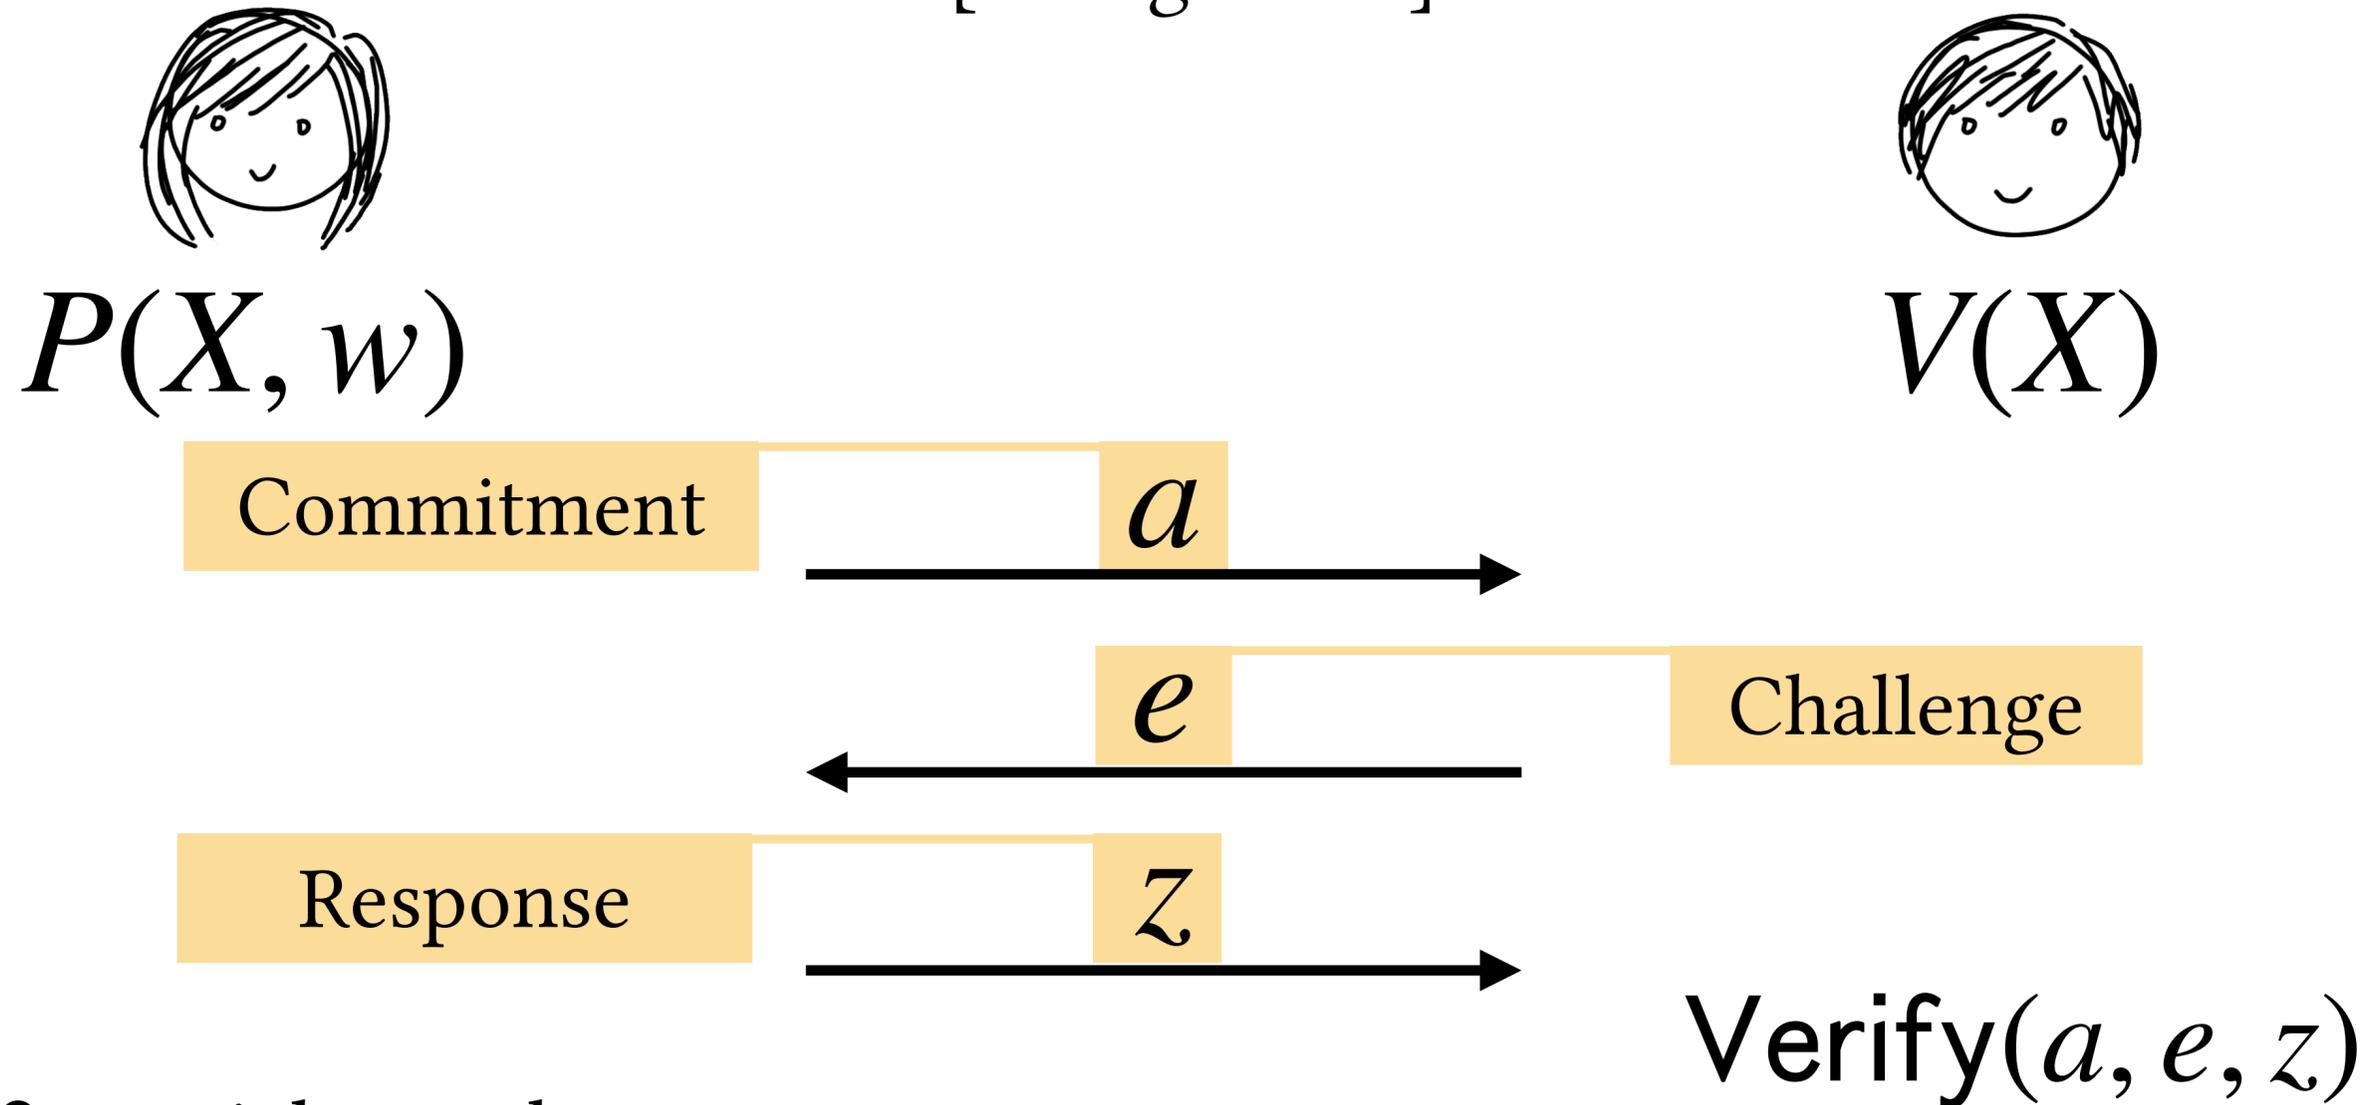


2-special soundness:

$$w \leftarrow \text{Ext}(X, a, (e_1, z_1), (e_2, z_2)) \text{ such that } R(X, w) = 1$$

Recap: Σ Protocol for Relation R

[Damgård 02]



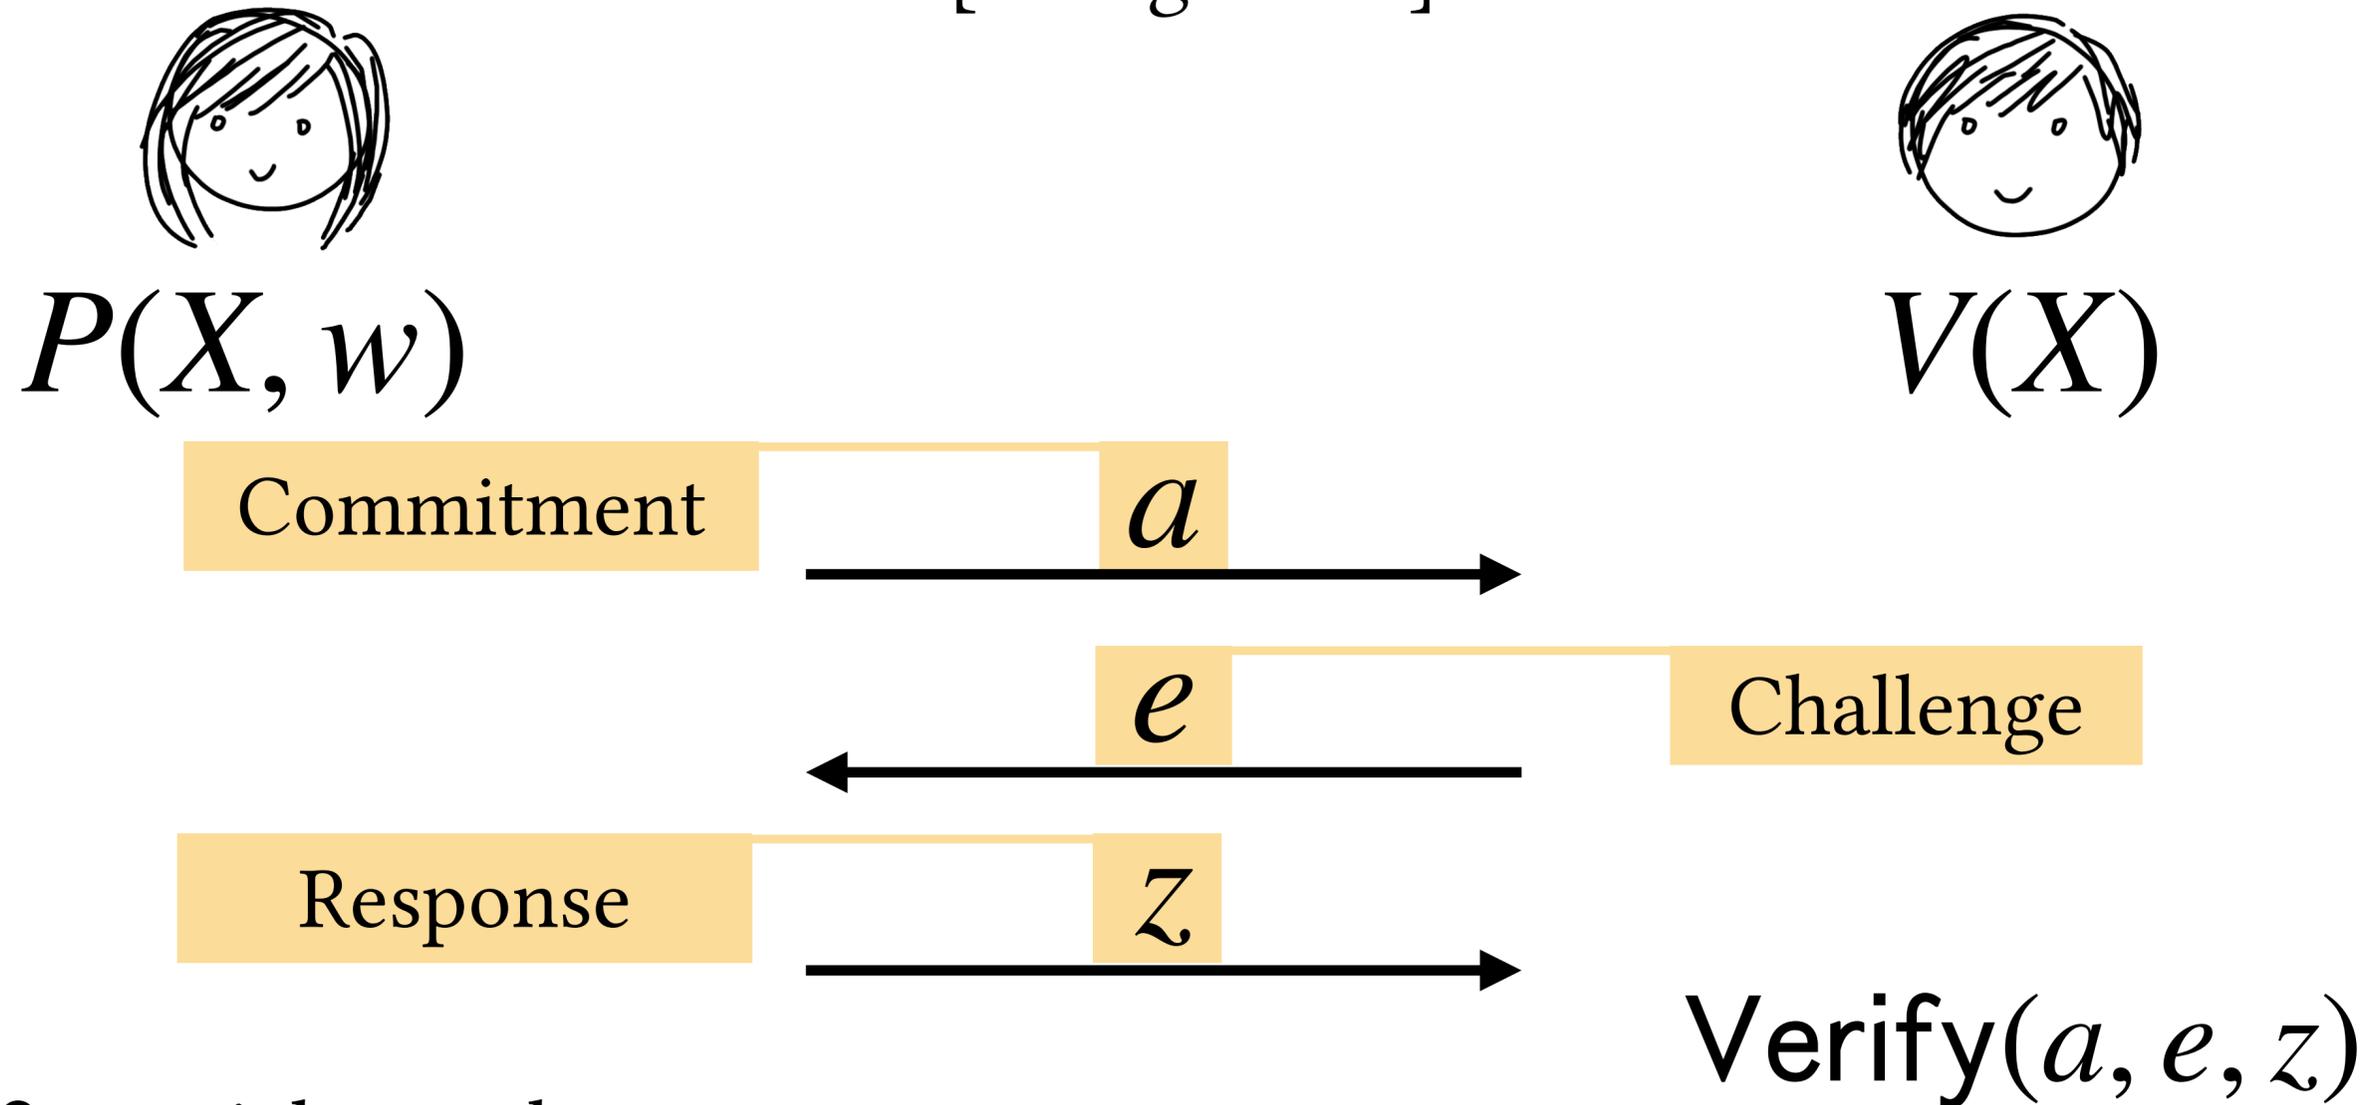
2-special soundness:

$w \leftarrow \text{Ext}(X, a, (e_1, z_1), (e_2, z_2))$ such that $R(X, w) = 1$

Fixed commitment

Recap: Σ Protocol for Relation R

[Damgård 02]



2-special soundness:

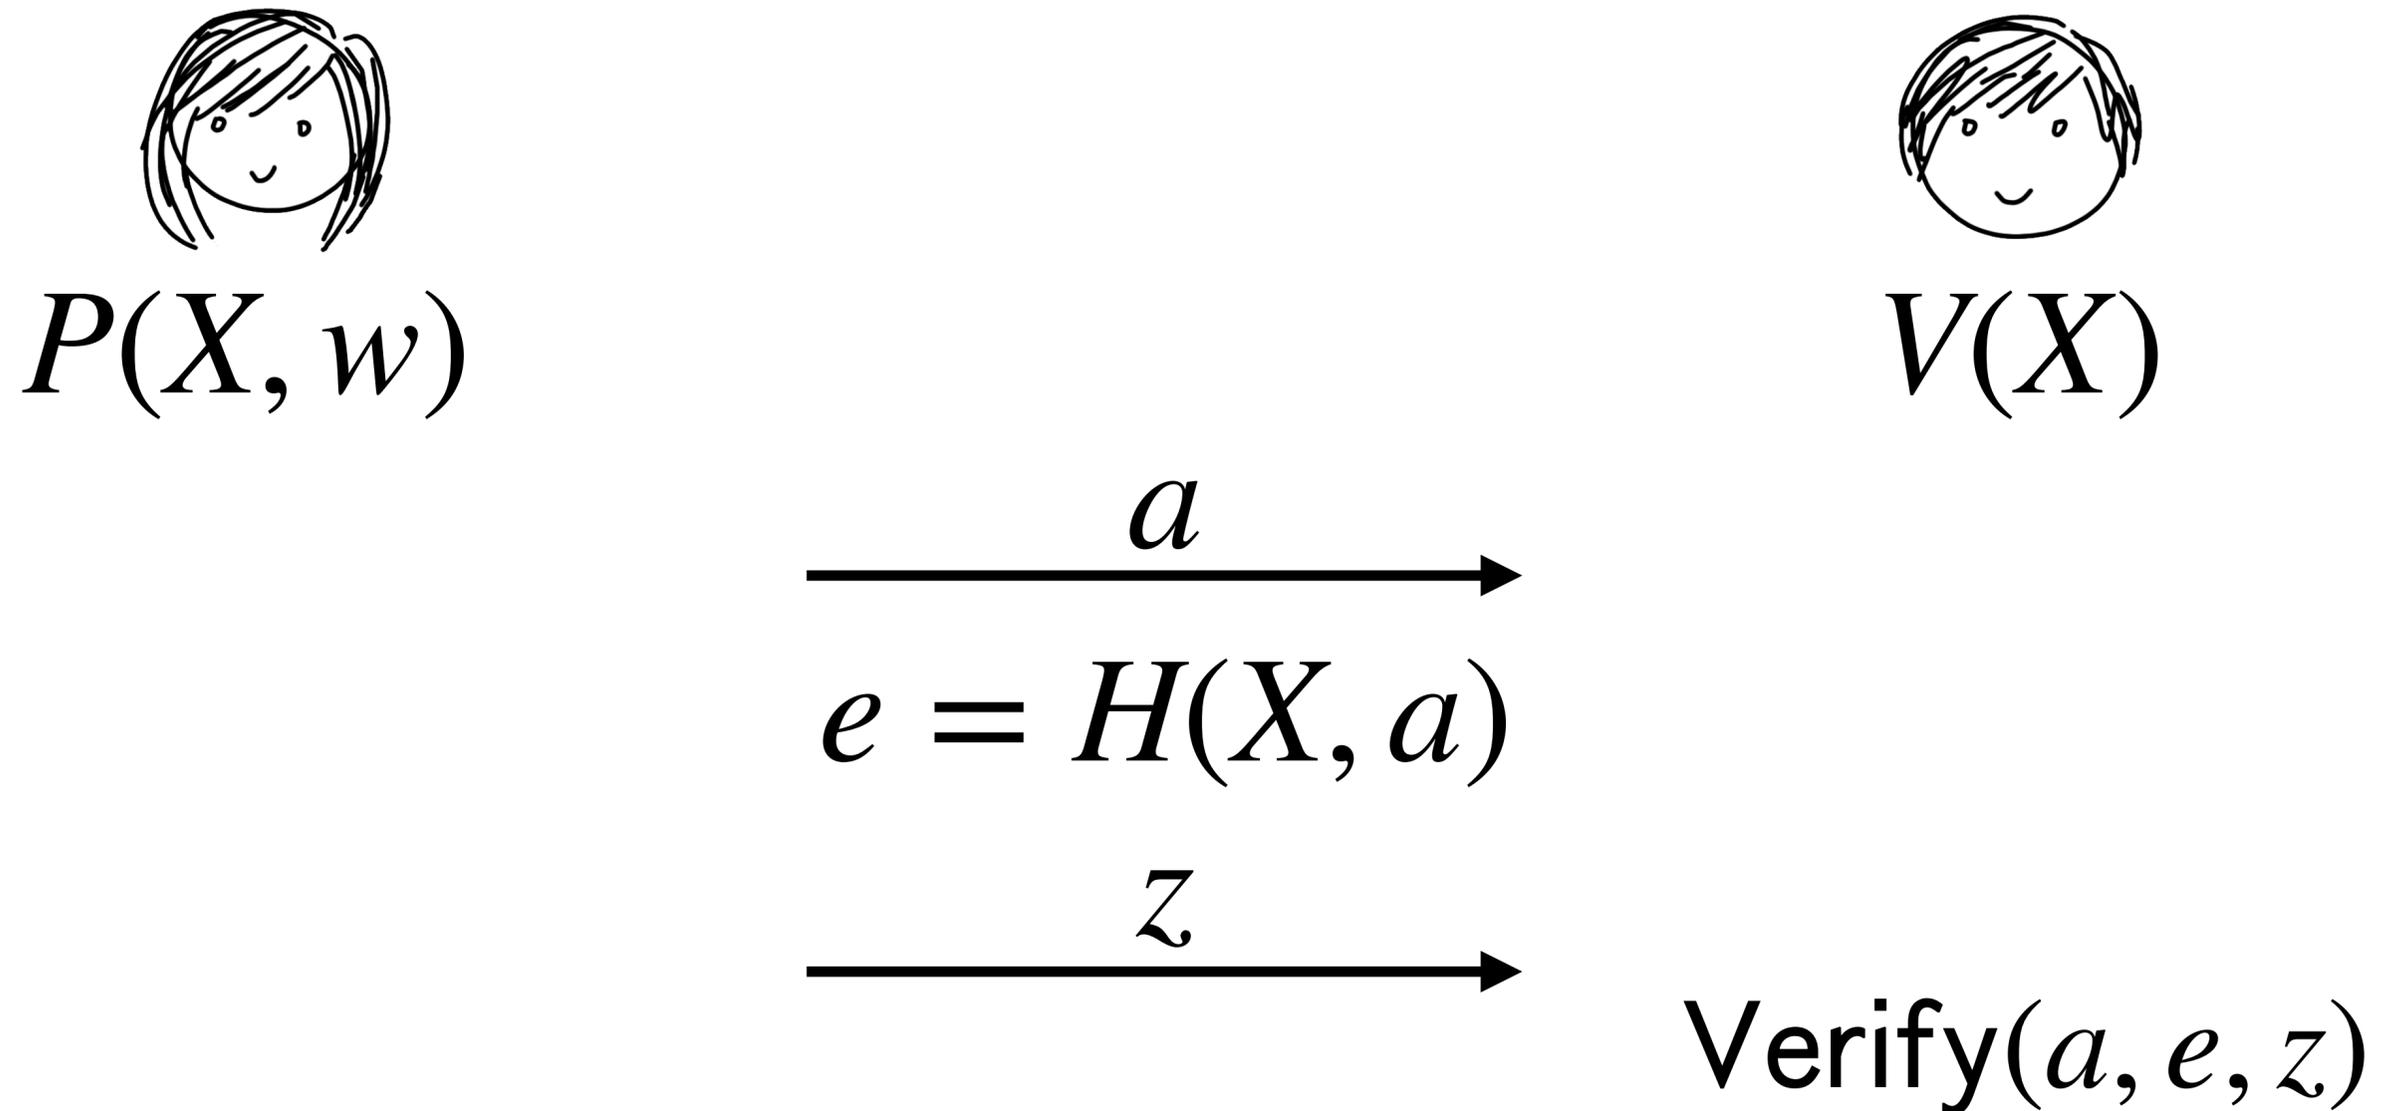
Varying (ch, resp) pairs

$w \leftarrow \text{Ext}(X, a, (e_1, z_1), (e_2, z_2))$ such that $R(X, w) = 1$

Fixed commitment

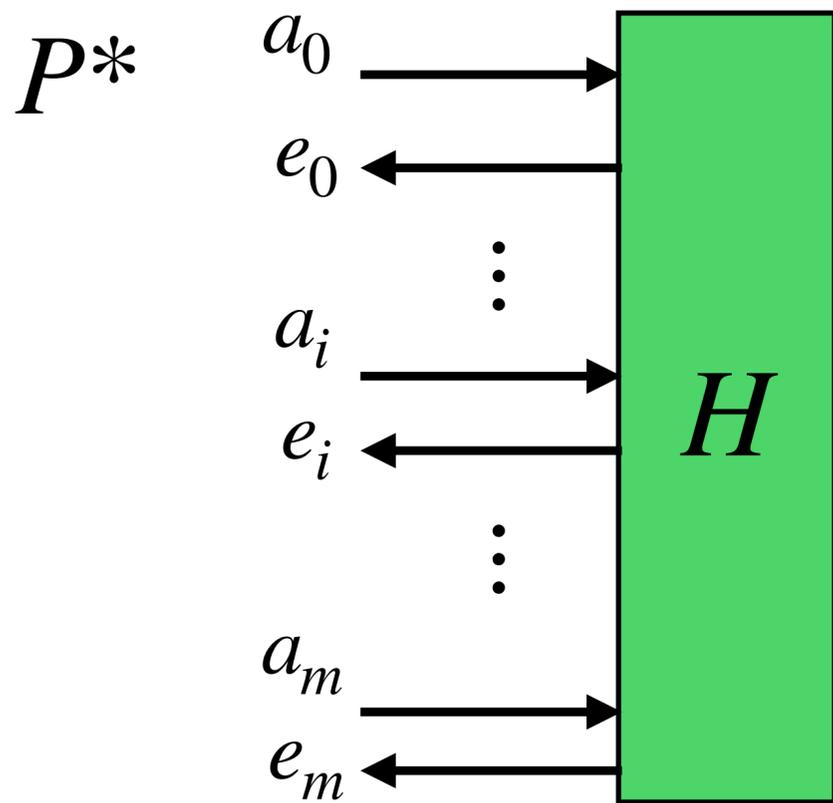
The Fiat-Shamir Transform

- [Fiat Shamir 87] provides a simple method to compile any public-coin protocol to a non-interactive proof, given a suitably chosen hash function



Fiat-Shamir: Security

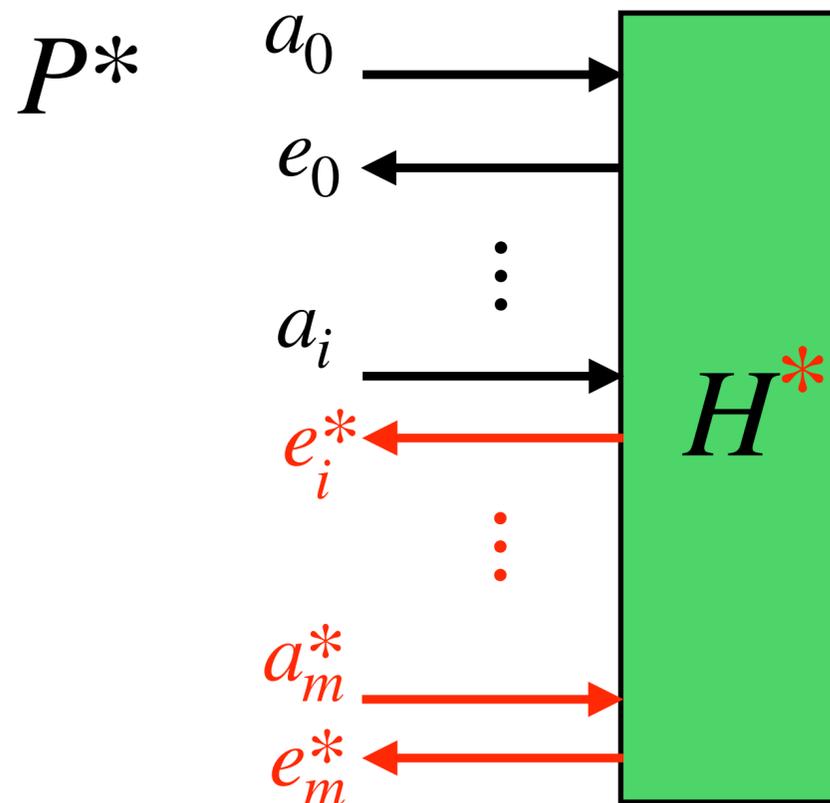
- “Forking” extraction strategy in Random Oracle Model [Pointcheval Stern 96]:



Output (a_i, e_i, z_i)

Probability of success:

p



Output (a_i, e_i^*, z_i^*)

p

$$\text{Ext} \left(\begin{matrix} (a_i, e_i) & (a_i, e_i^*) \\ z_i, z_i^* \end{matrix} \right)$$

Outputs witness w

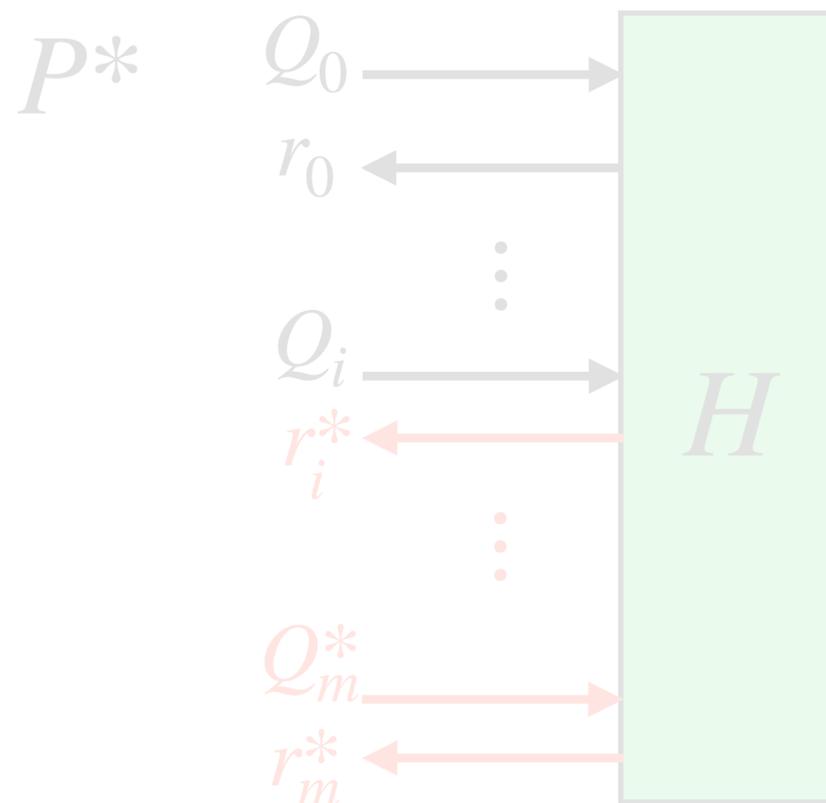
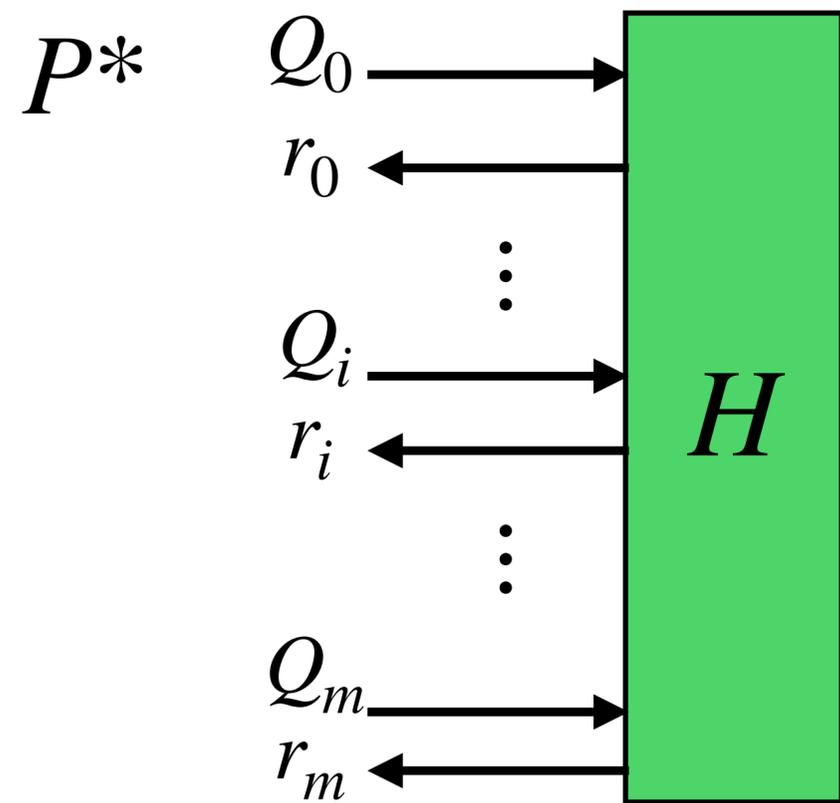
$\approx p^2$

Fiat-Shamir Compilation

- Advantages:
 - Simple to describe/implement
 - Very efficient; proving, verification cost exactly the same as input Σ -protocol
- Downsides:
 - Forking strategy does not compose;
unclear how to prove concurrent security
 - Quadratic security loss

Straight-line Extraction

- Formalized by [Pass 03] in the Random Oracle Model:



$\text{Ext}((Q_0, r_0), \dots, (Q_m, r_m))$

Outputs witness w

Probability of success:

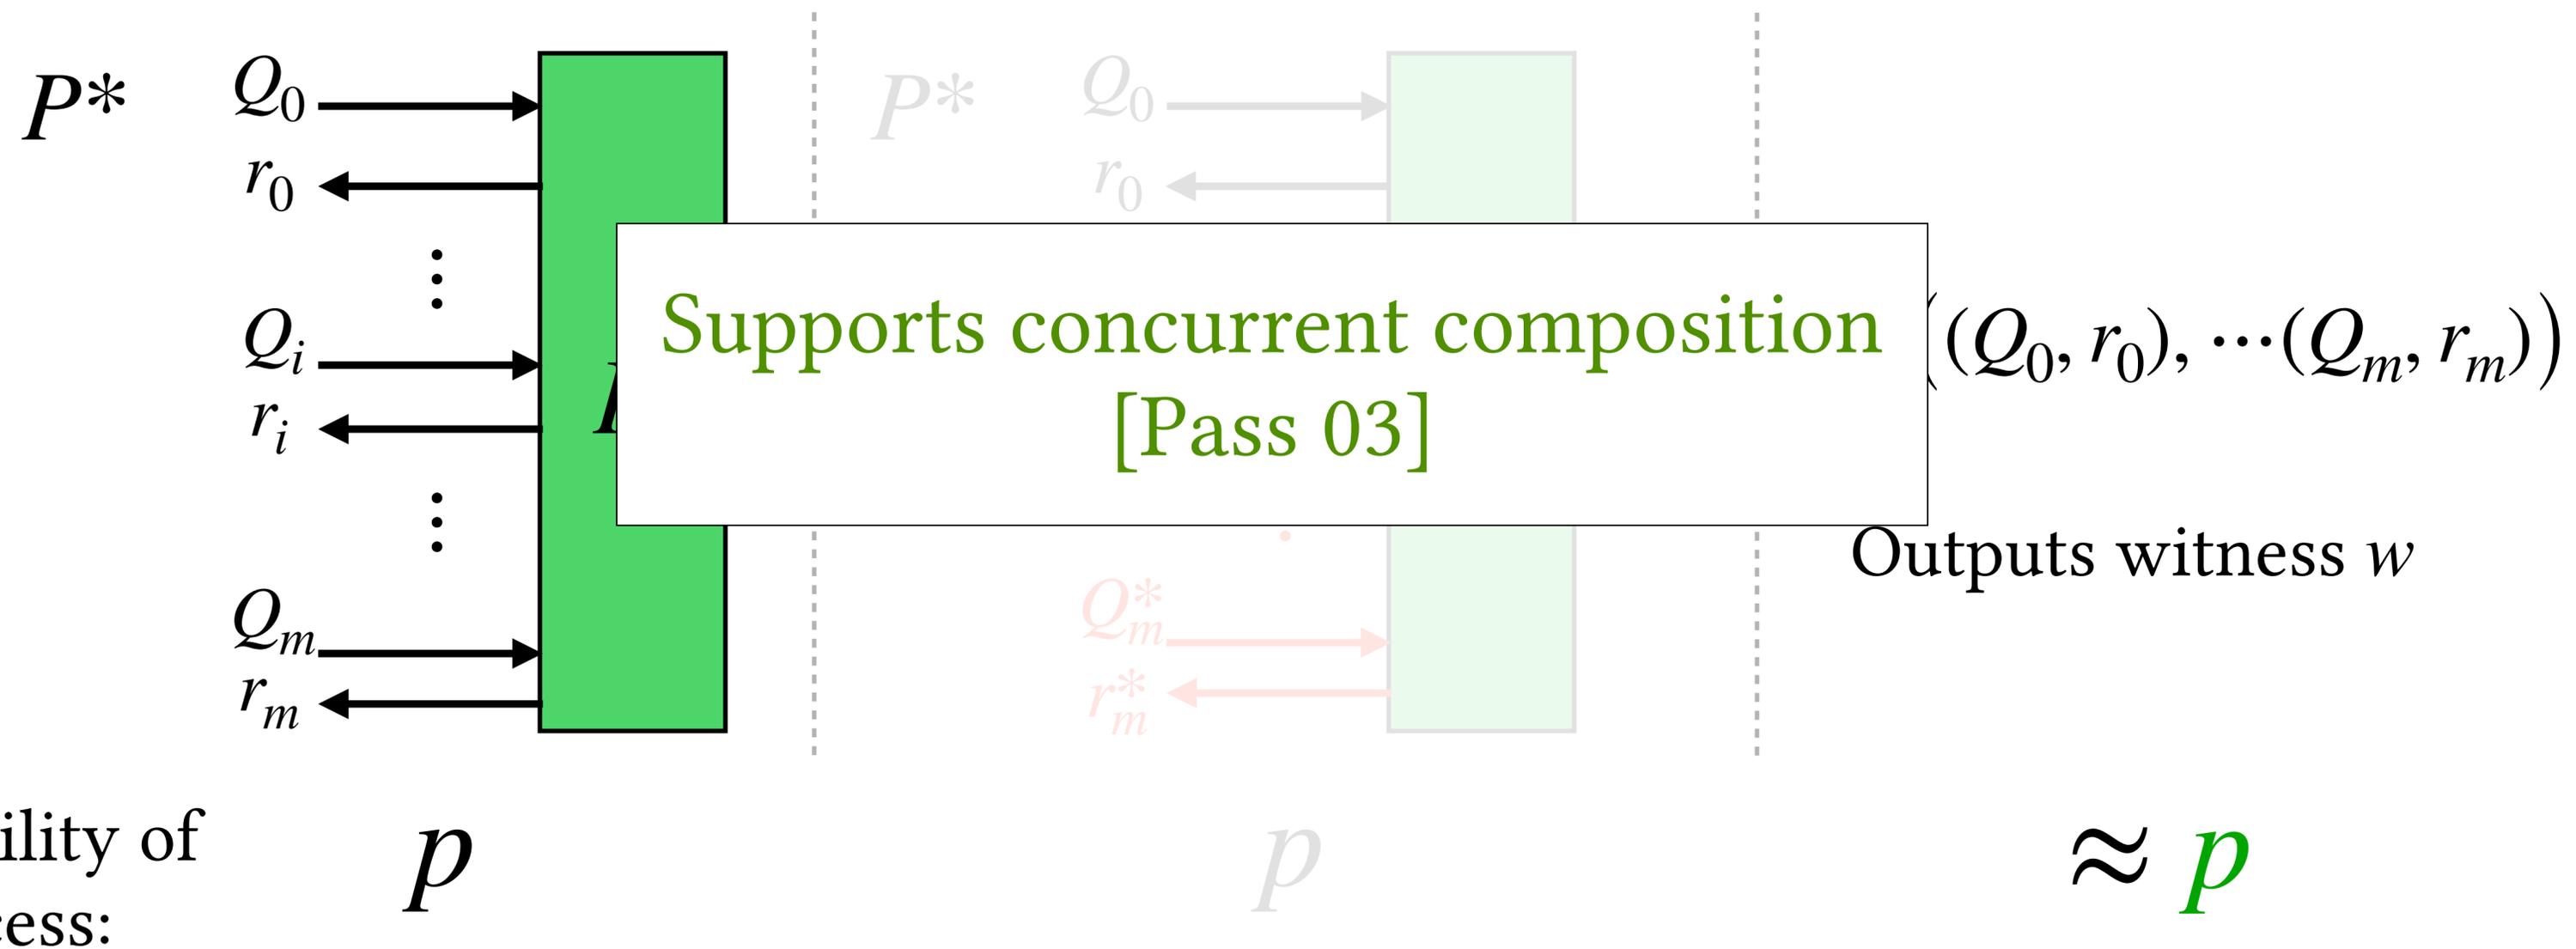
p

p

$\approx p$

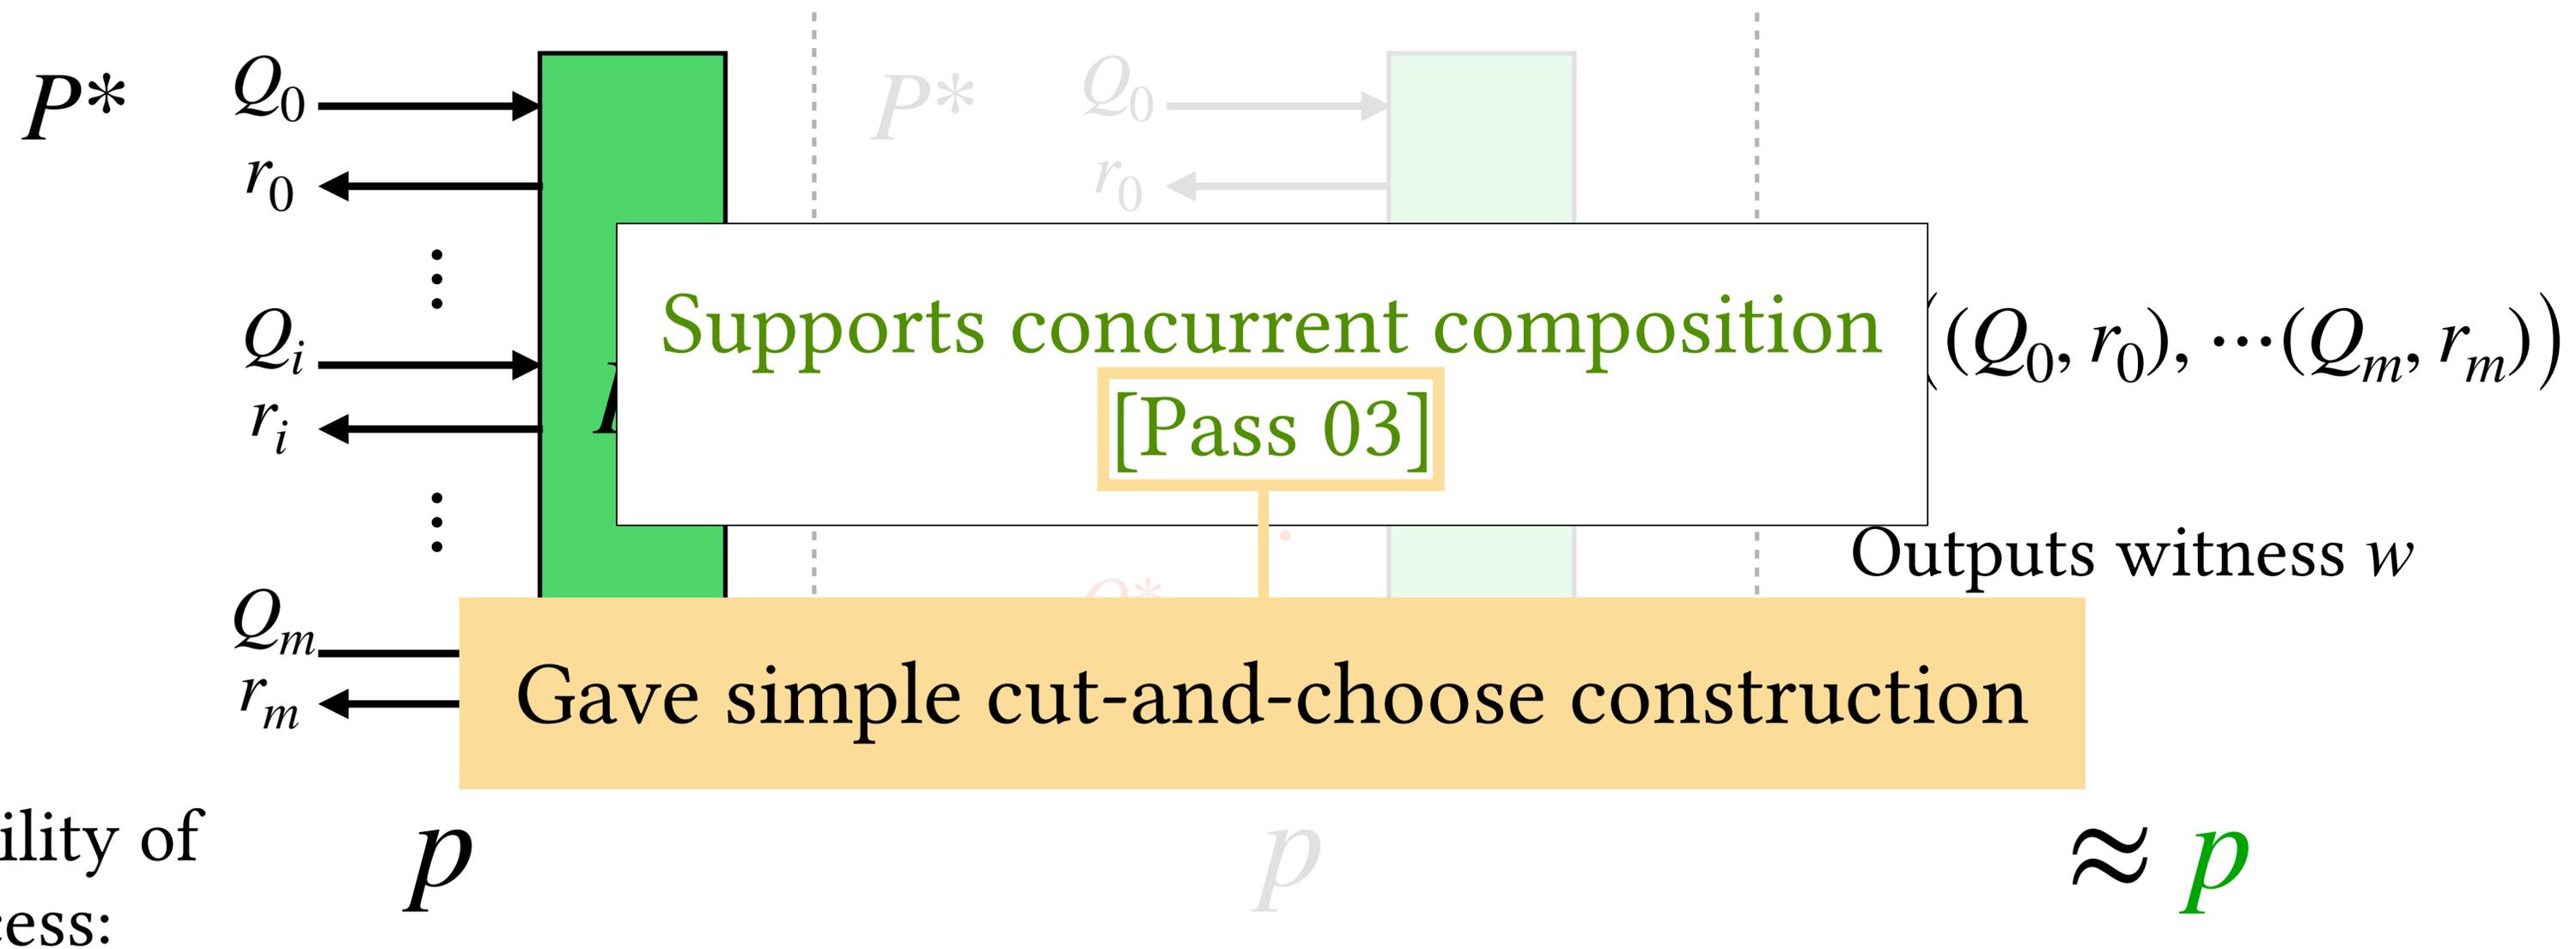
Straight-line Extraction

- Formalized by [Pass 03] in the Random Oracle Model:



Straight-line Extraction

- Formalized by [Pass 03] in the Random Oracle Model:



Fischlin's Transformation

- [Fischlin 05] gave a straight-line extractable compiler that avoids cut-and-choose logistics through a clever “proof of work” type idea

$P(X, w)$

$\xrightarrow{a, e, z}$

$V(X)$

$H(a, e, z) \stackrel{?}{=} 0$

Verify(a, e, z)

Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

$P(X, w) :$



Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

$P(X, w)$: Sample Σ -protocol first message ' a '



Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

$P(X, w)$: Sample Σ -protocol first message ' a '

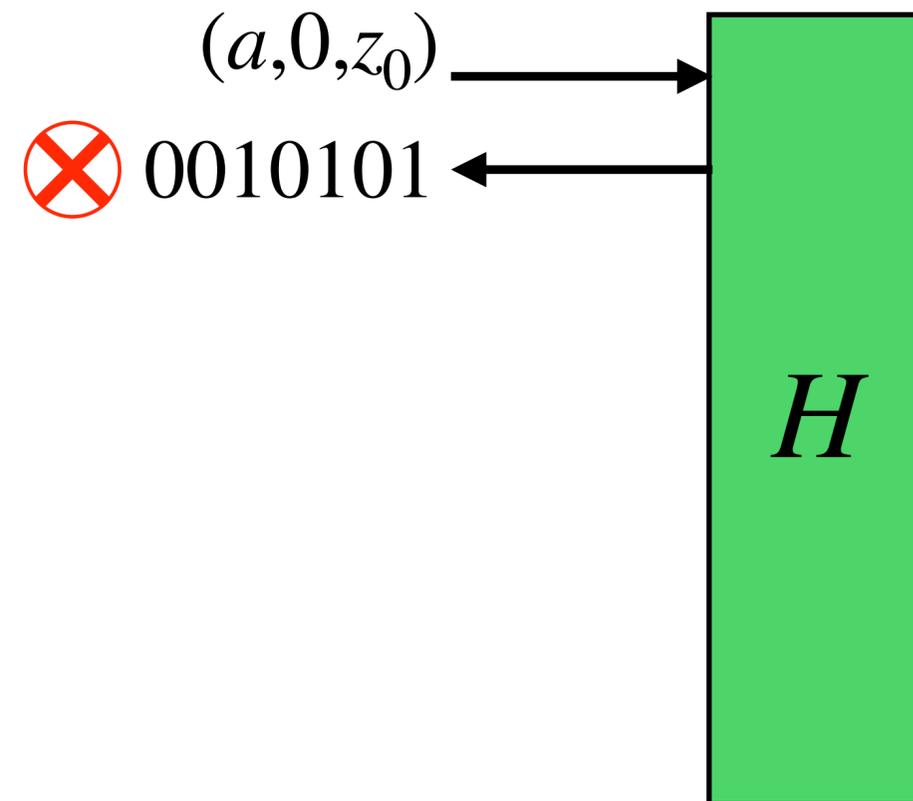
$(a, 0, z_0)$



Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

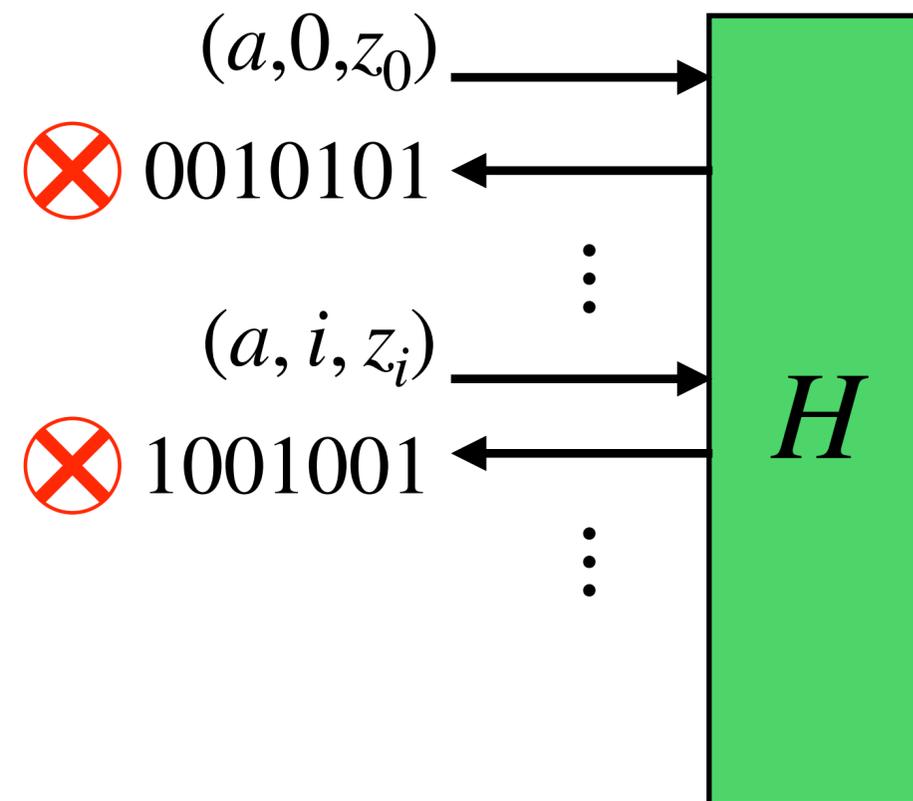
$P(X, w)$: Sample Σ -protocol first message ' a '



Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

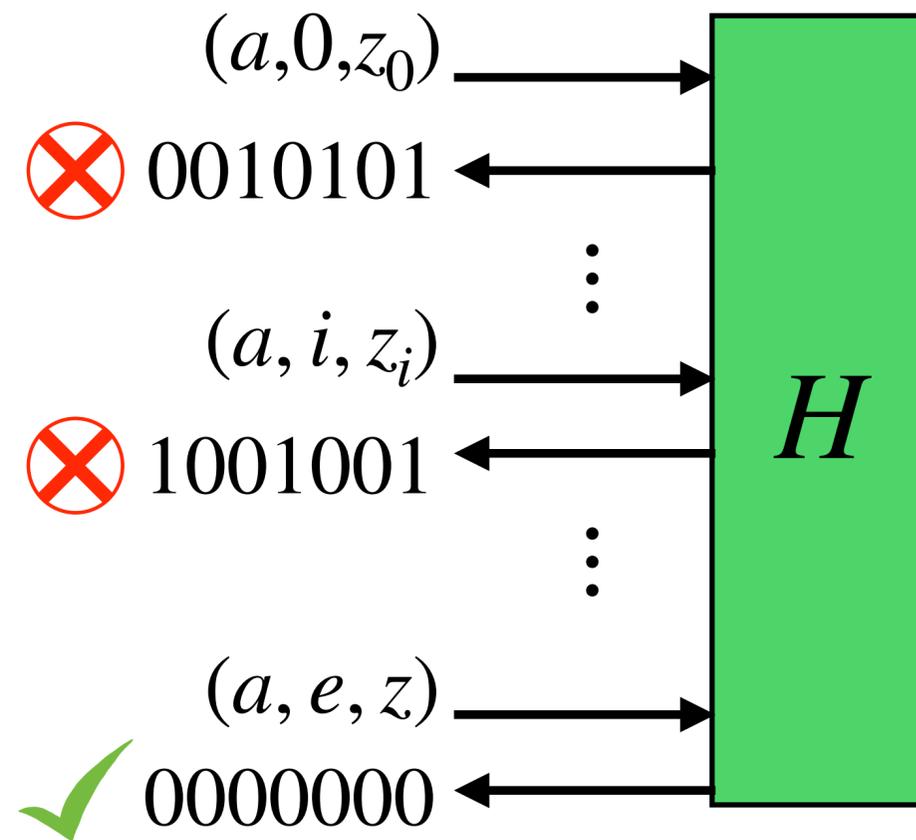
$P(X, w)$: Sample Σ -protocol first message ' a '



Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

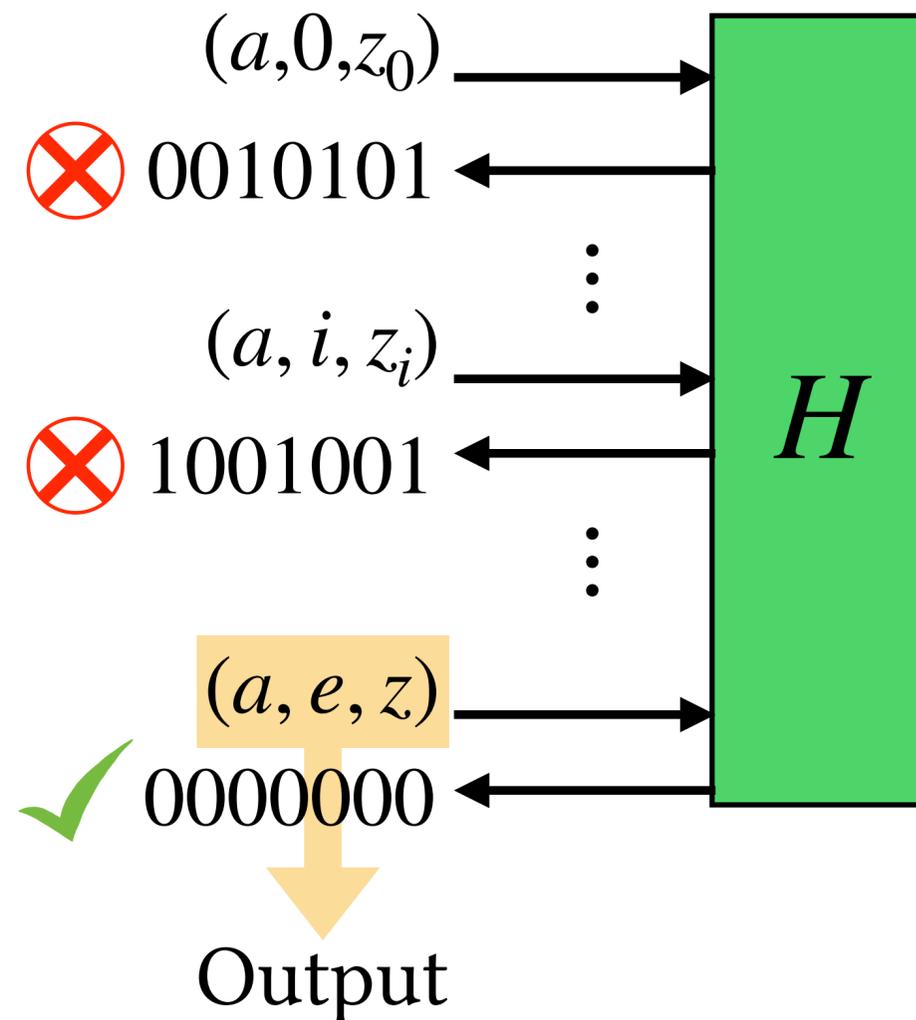
$P(X, w)$: Sample Σ -protocol first message 'a'



Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

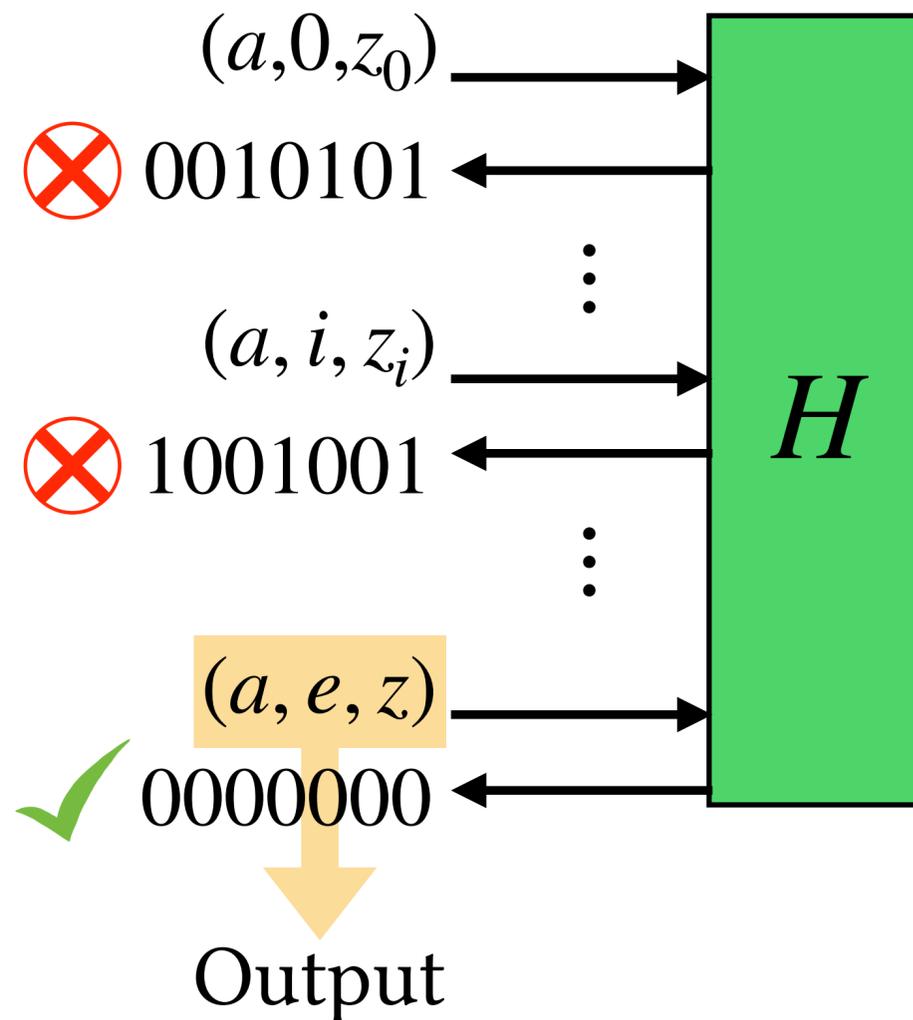
$P(X, w)$: Sample Σ -protocol first message 'a'



Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

$P(X, w)$: Sample Σ -protocol first message 'a'



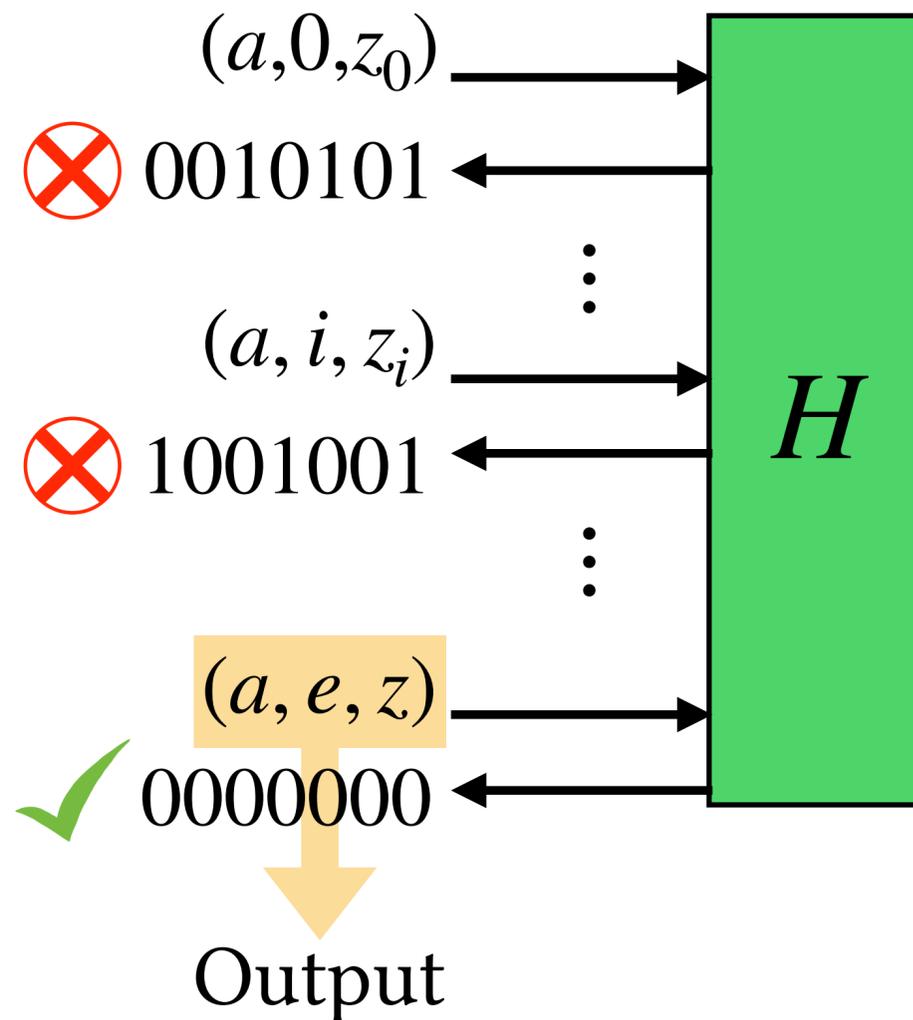
Soundness: Except with $\Pr=2^{-\ell}$, P is forced to query more than one accepting transcript to H

Completeness: P terminates in poly time when ℓ is small, i.e. $O(\log \kappa)$

Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

$P(X, w)$: Sample Σ -protocol first message 'a'



Soundness: Except with $\Pr=2^{-\ell}$, P is forced to query more than one accepting transcript to H

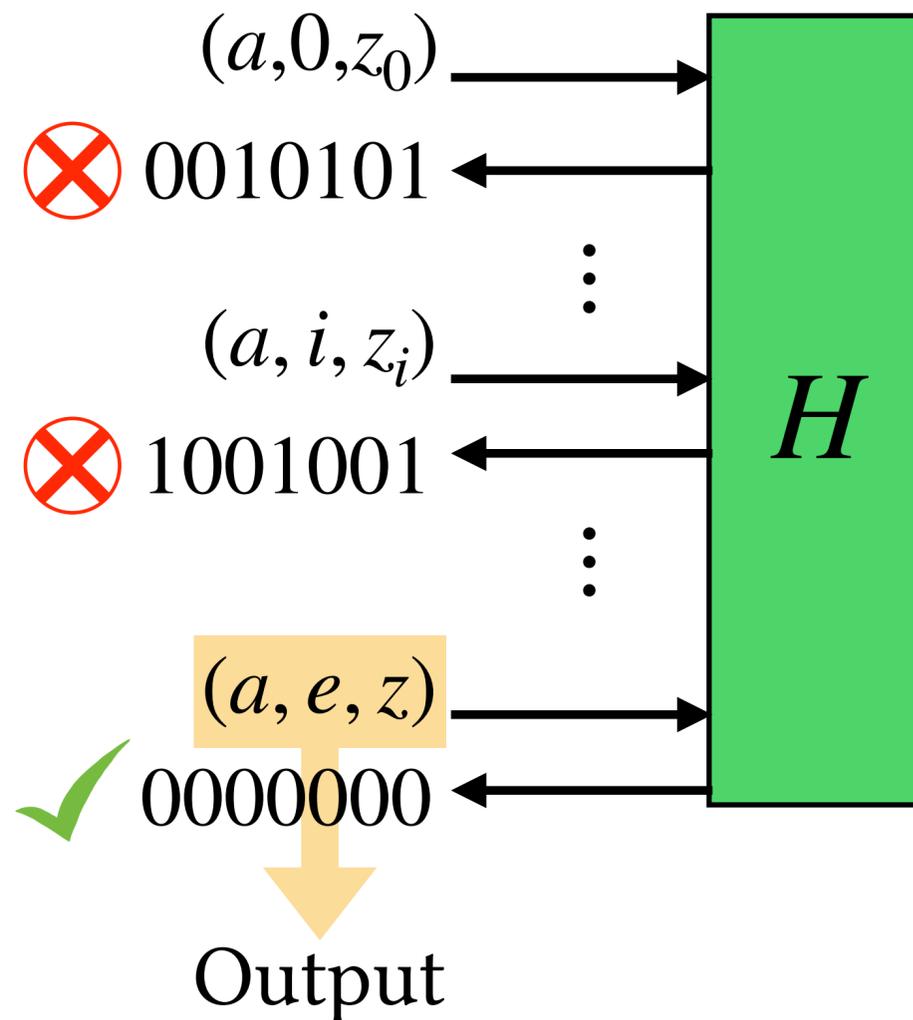
Completeness: P terminates in poly time when ℓ is small, i.e. $O(\log \kappa)$

Problem!

Fischlin's Transformation

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

$P(X, w)$: Sample Σ -protocol first message 'a'



Soundness: Except with $\Pr=2^{-\ell}$, P is forced to query more than one accepting transcript to H

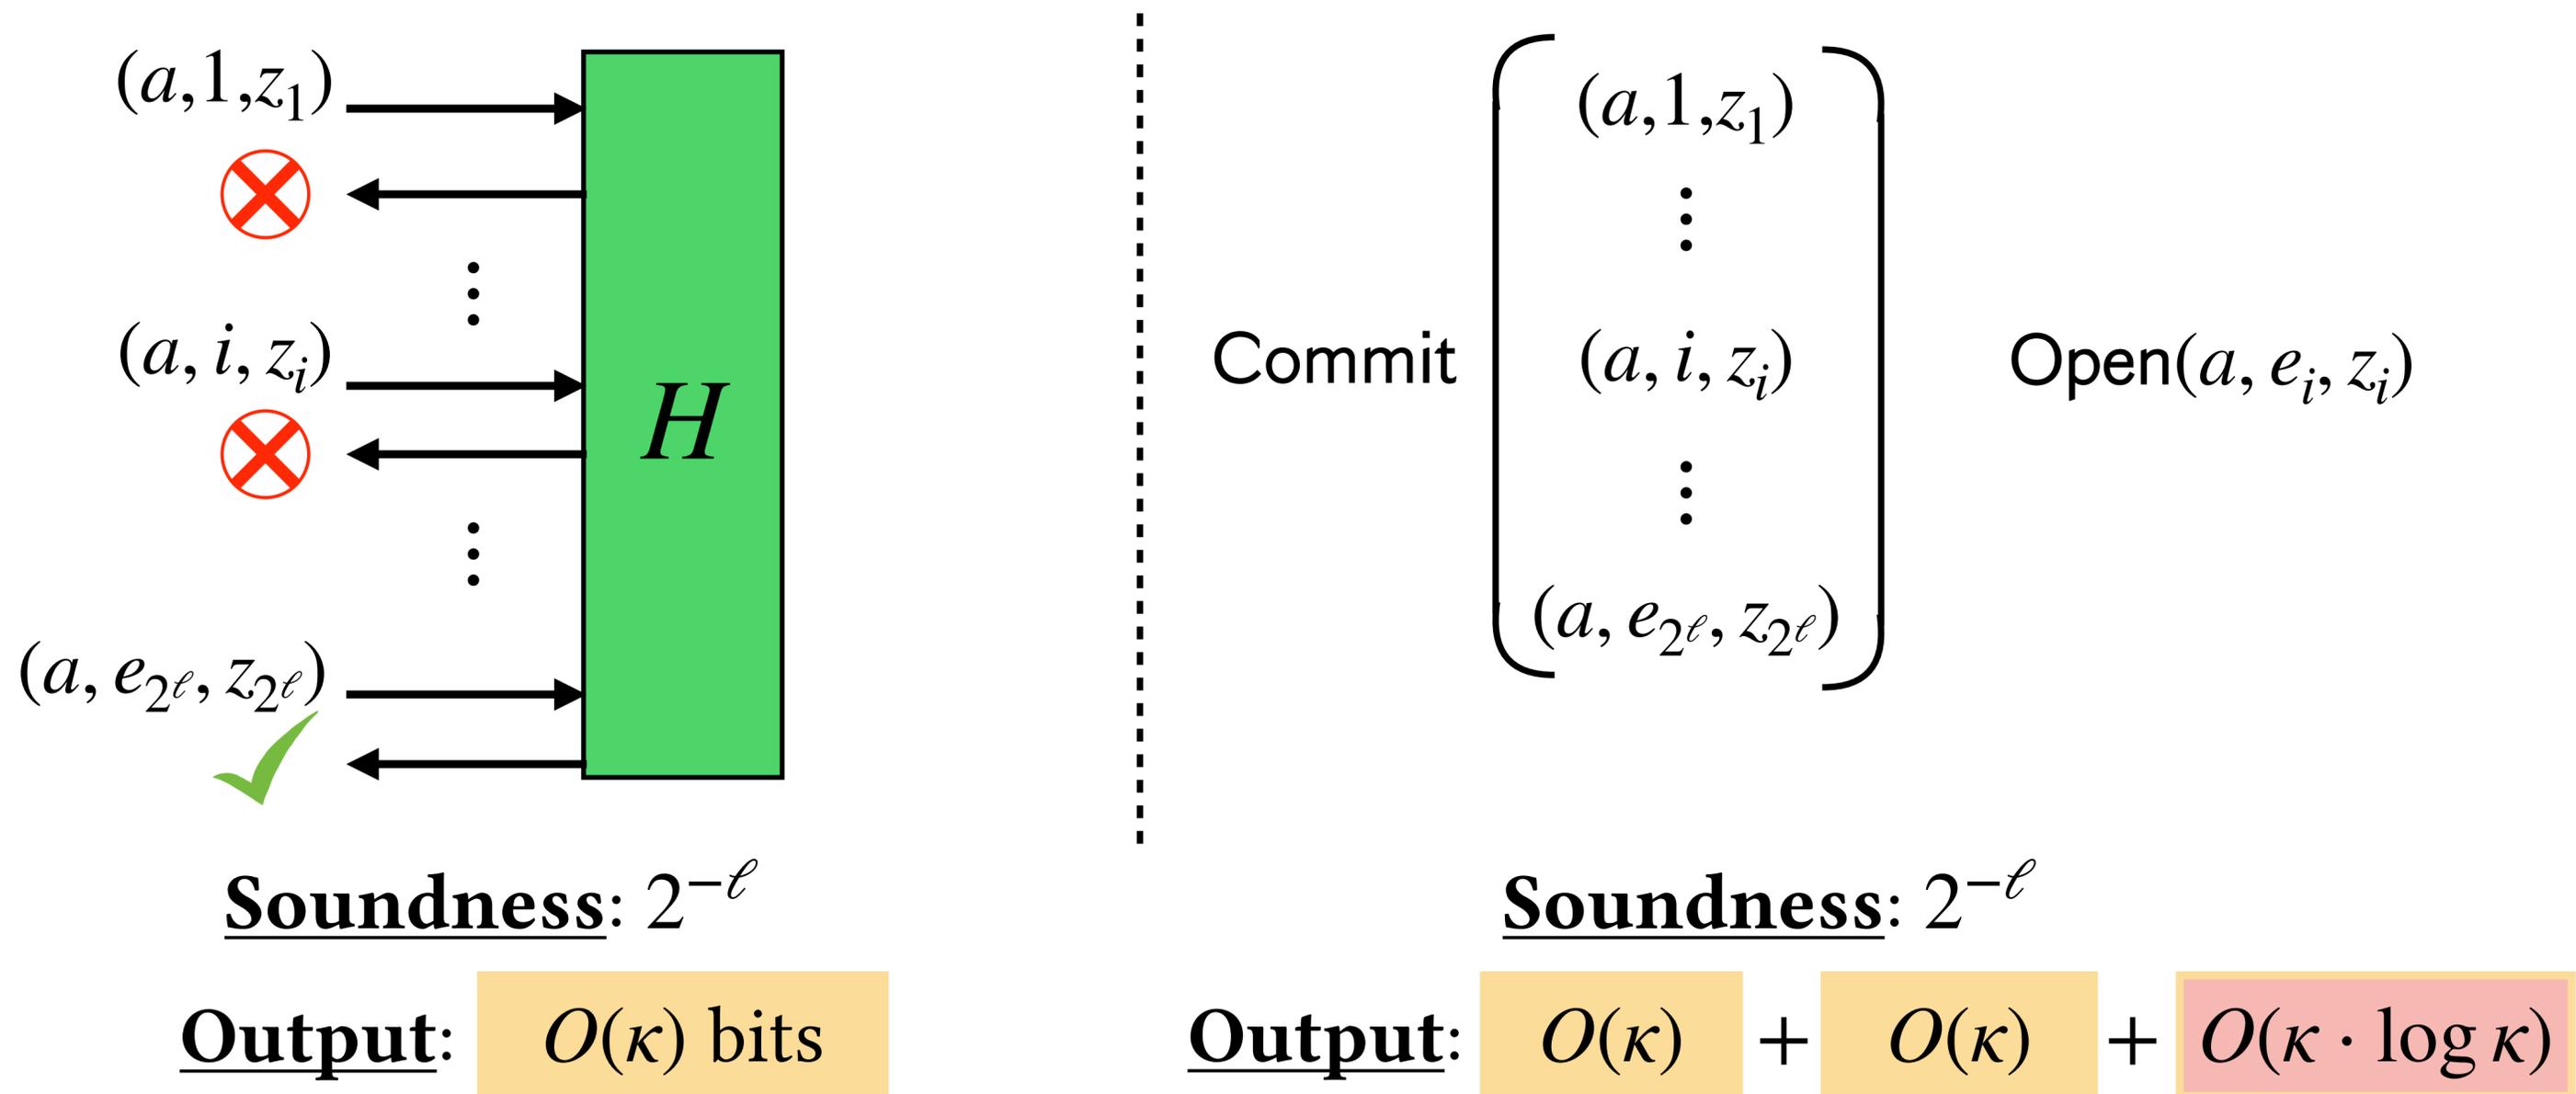
Completeness: P terminates in poly time when ℓ is small, i.e. $O(\log \kappa)$

Problem!

Full Soundness: Repeat r times

Fischlin05 vs Pass03

$P(X, w)$: Sample Σ -protocol first message 'a'



Fischlin05 vs Pass03: Qualitative

- Pass' compiler works for *any* Sigma protocol
- Fischlin's compiler works for a restricted class of Sigma protocols with 'quasi-unique responses'
- Supported by many standard Sigma protocols (eg. DLog), but many *may* not—especially if a statement can have multiple witnesses (eg. Pedersen Commitment opening, 1-of-2 witnesses, etc.)

Quasi-unique Responses [Fischlin 05]

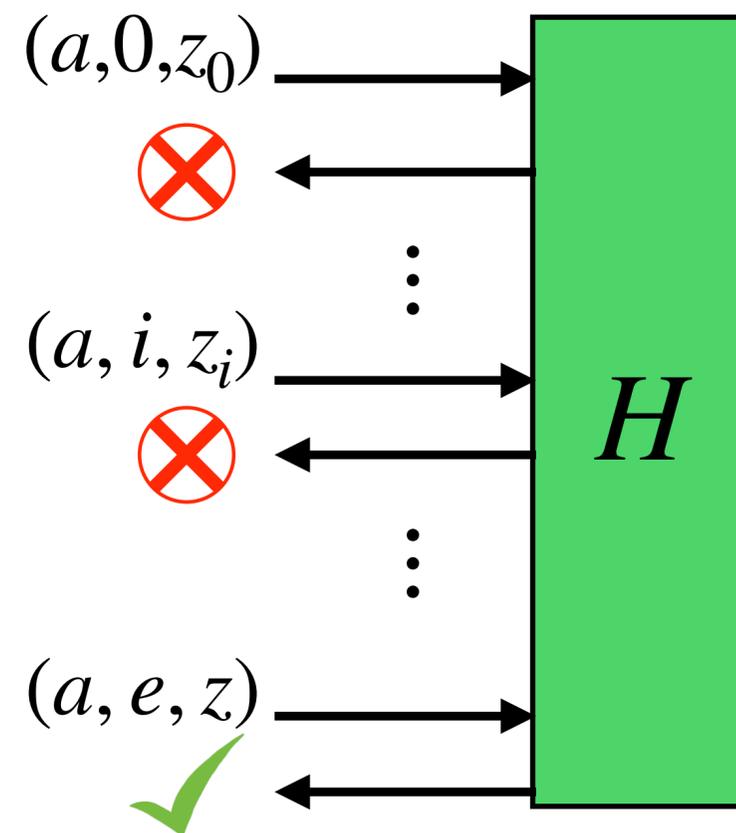
Hard: $(a, e, z, z') \leftarrow \mathcal{A}(\text{pp})$ such that
 $V(a, e, z) = V(a, e, z') = 1$

Fixing (a, e) fixes z

Quasi-unique Responses [Fischlin 05]

Hard: $(a, e, z, z') \leftarrow \mathcal{A}(\text{pp})$ such that
 $V(a, e, z) = V(a, e, z') = 1$

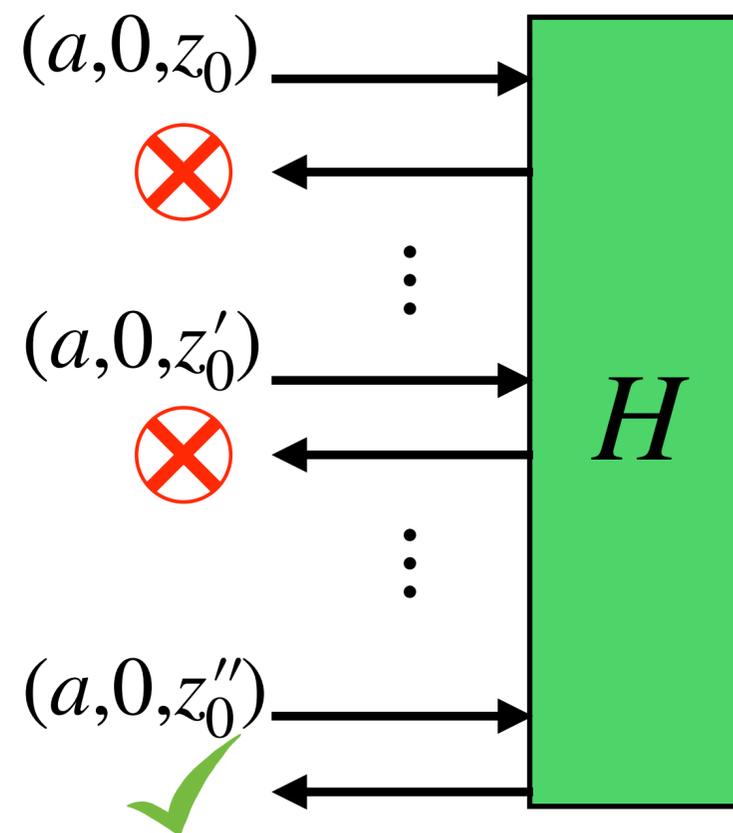
Fixing (a, e) fixes z



Quasi-unique Responses [Fischlin 05]

Hard: $(a, e, z, z') \leftarrow \mathcal{A}(\text{pp})$ such that
 $V(a, e, z) = V(a, e, z') = 1$

Fixing (a, e) fixes z

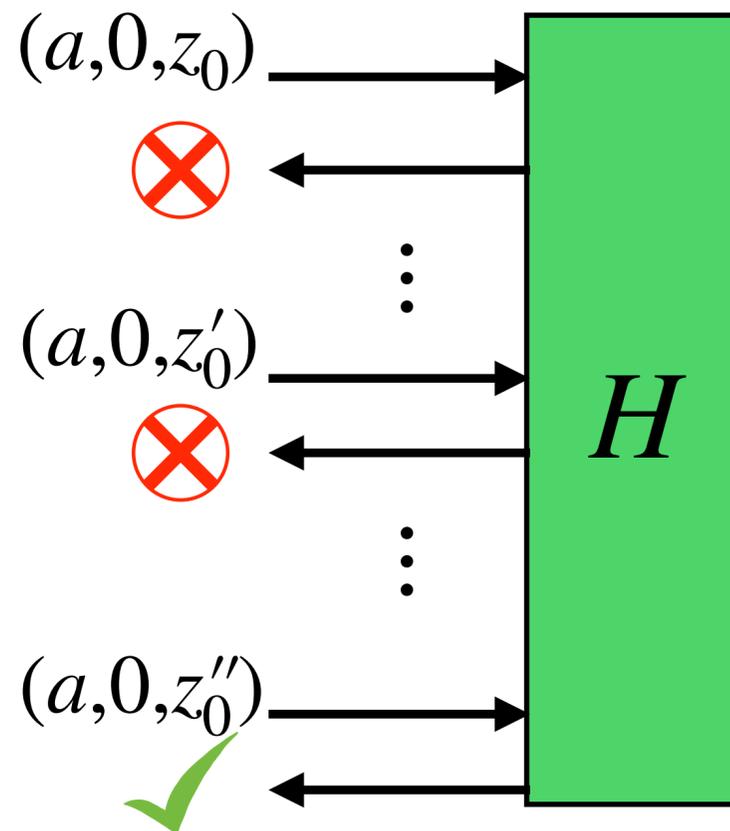


Easy to see how this
ties into soundness of
Fischlin's compiler

Quasi-unique Responses [Fischlin 05]

Hard: $(a, e, z, z') \leftarrow \mathcal{A}(\text{pp})$ such that
 $V(a, e, z) = V(a, e, z') = 1$

Fixing (a, e) fixes z



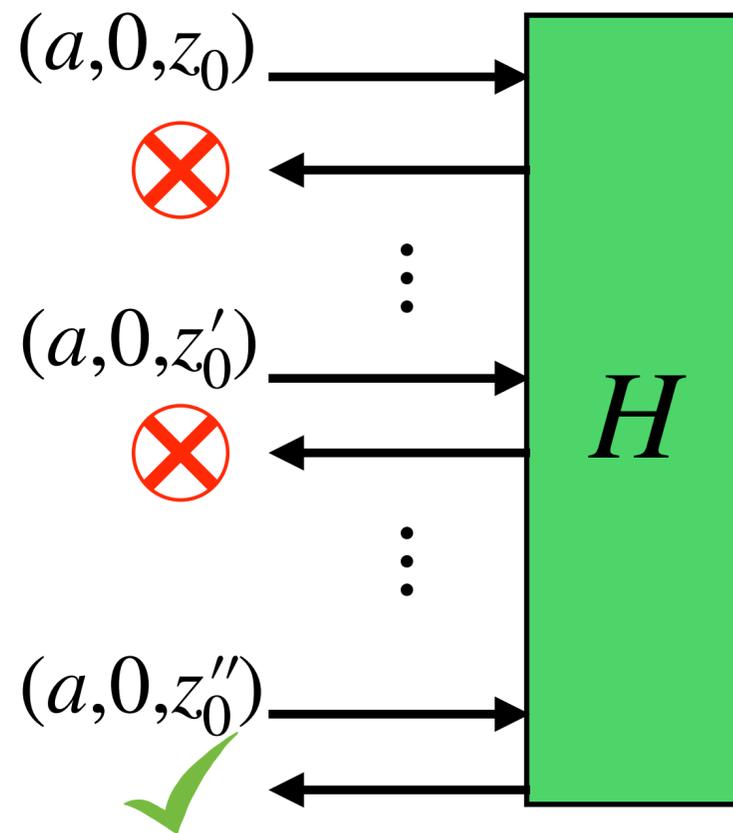
Easy to see how this
ties into soundness of
Fischlin's compiler

Prover can produce a proof
without ever having to try
more than one challenge

Quasi-unique Responses [Fischlin 05]

Hard: $(a, e, z, z') \leftarrow \mathcal{A}(\text{pp})$ such that
 $V(a, e, z) = V(a, e, z') = 1$

Fixing (a, e) fixes z



Easy to see how this
ties into soundness of
Fischlin's compiler

Prover can produce a proof
without ever having to try
more than one challenge

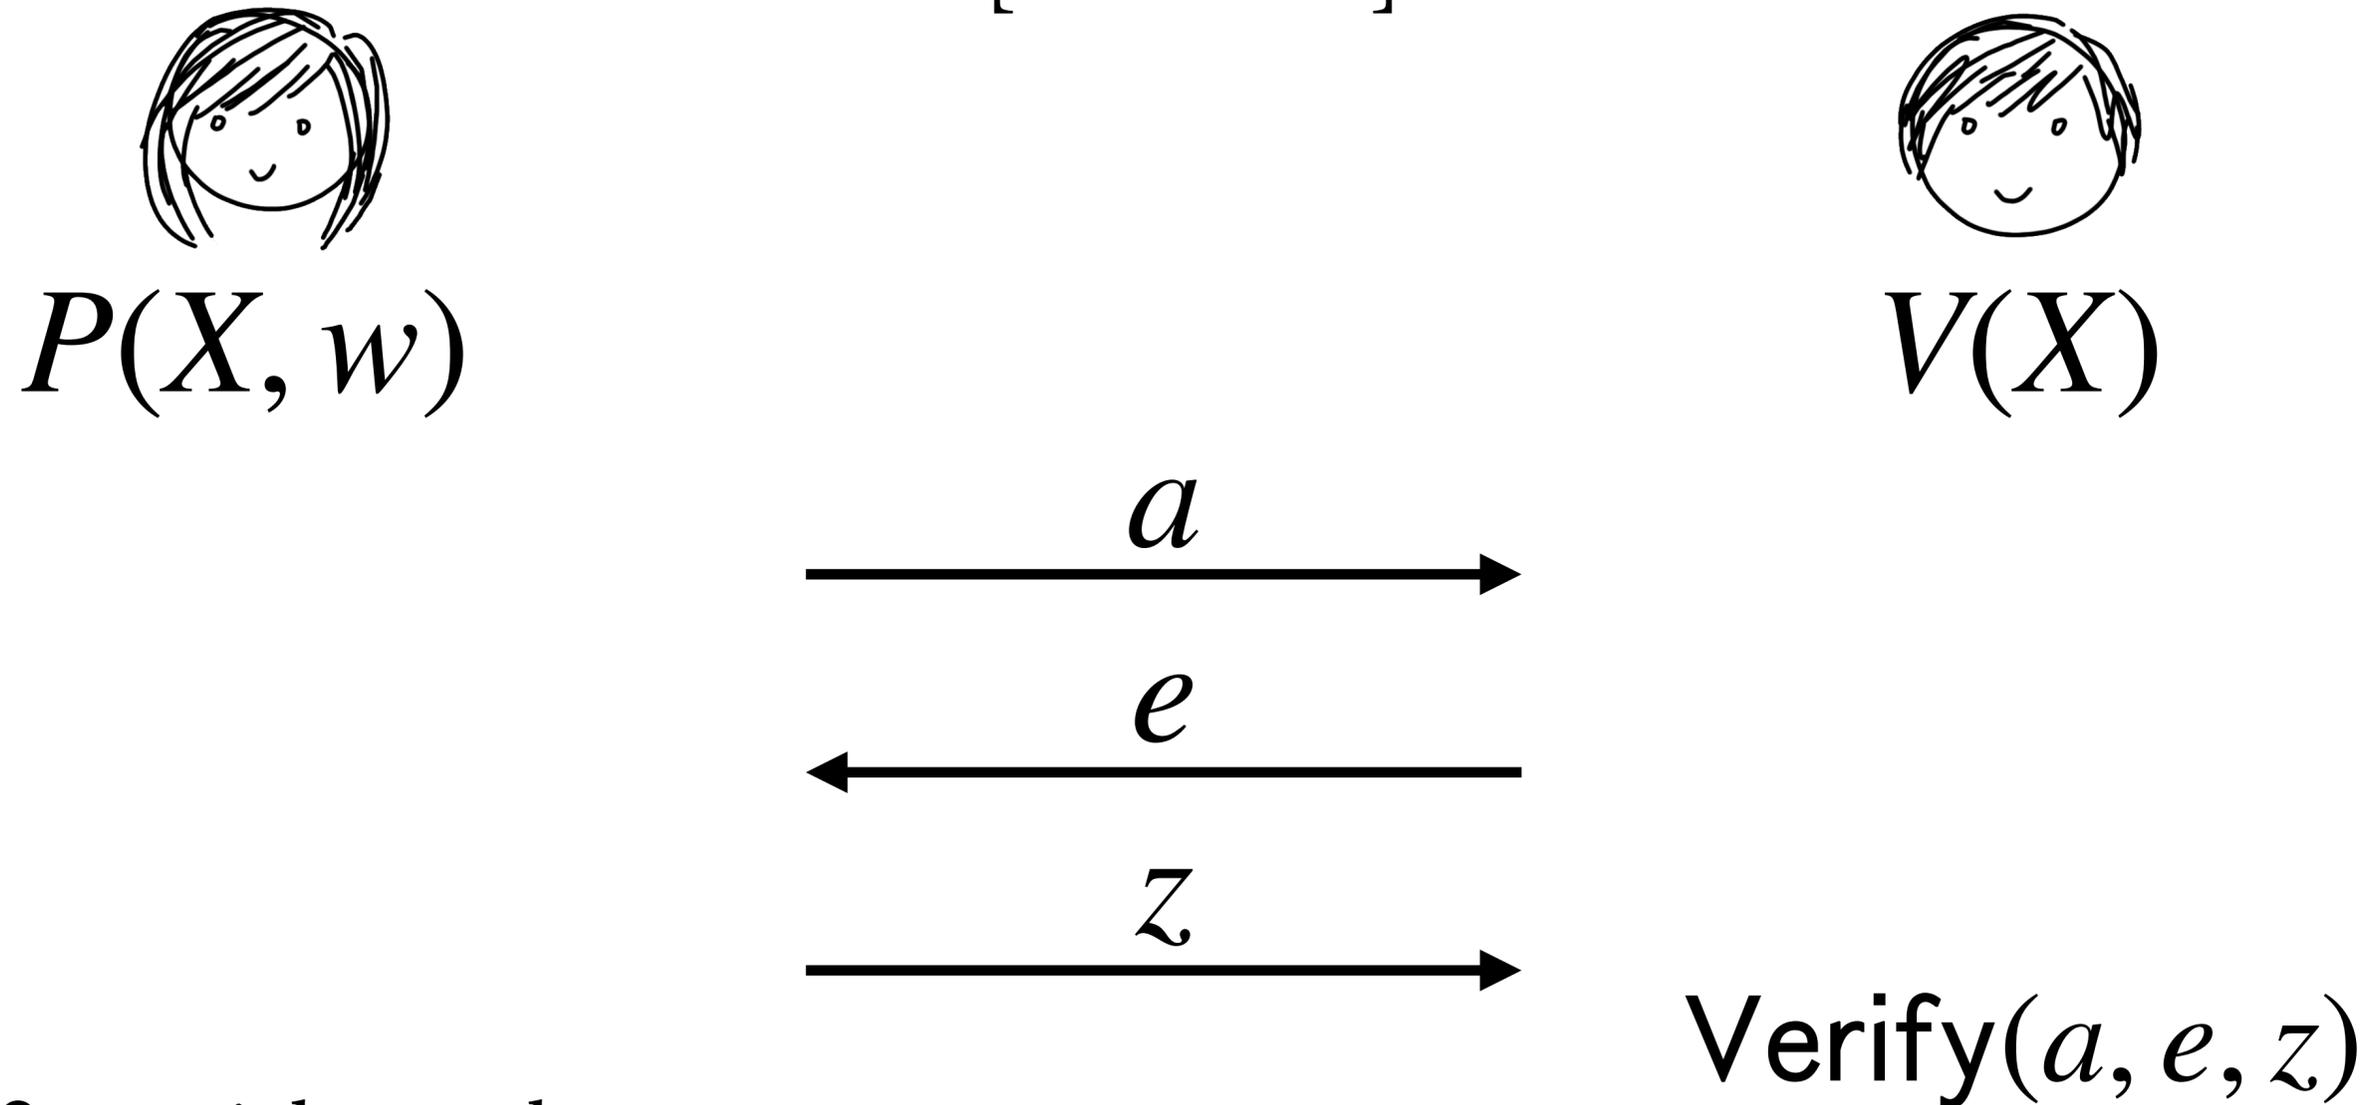
Recall:
Extractor needs transcripts
with different challenges

Is it *really* necessary, though?

- Folklore: breaking Sigma protocol abstraction, and simply ‘adjusting syntax’ of the extractor is usually sufficient to preserve Proof of Knowledge
- This is demonstrated by the Sigma protocol to prove knowledge of one-out-of-two witnesses
[Cramer Damgård Schoenmakers 94]
- Intuition: (a, e, z, z') allow for the extraction of a witness

Tightening Conditions for Extraction

[This work]



2-special soundness:

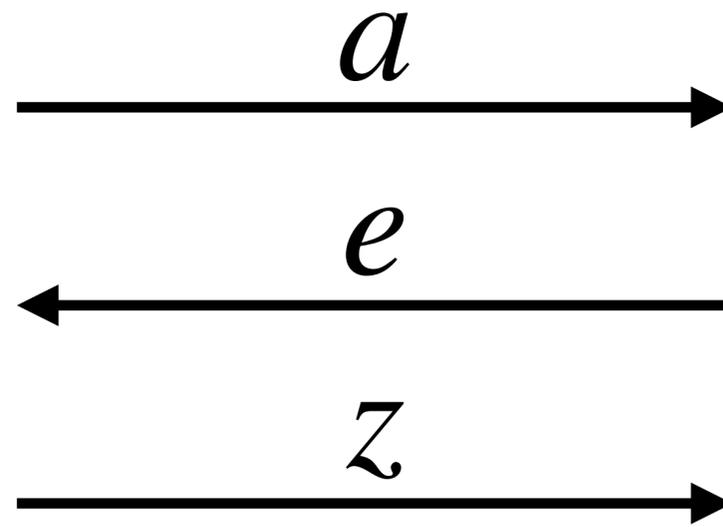
$$w \leftarrow \text{Ext}(X, a, (e_1, z_1), (e_2, z_2)) \text{ such that } R(X, w) = 1$$

Tightening Conditions for Extraction

[This work]


 $P(X, w)$


 $V(X)$



$\text{Verify}(a, e, z)$

Strong 2-special soundness:

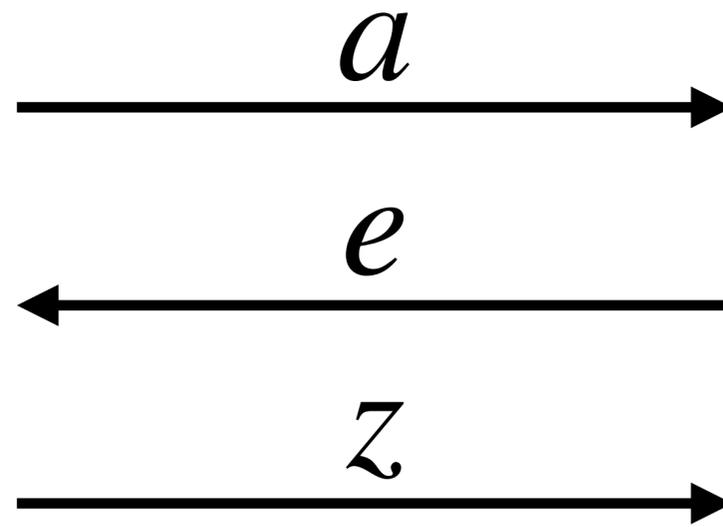
$w \leftarrow \text{Ext}(X, a, (e_1, z_1), (e_2, z_2))$ such that $R(X, w) = 1$

Tightening Conditions for Extraction

[This work]


 $P(X, w)$


 $V(X)$



$\text{Verify}(a, e, z)$

Strong 2-special soundness:

$w \leftarrow \text{Ext}(X, a, (e_1, z_1), (e_2, z_2))$ such that $R(X, w) = 1$

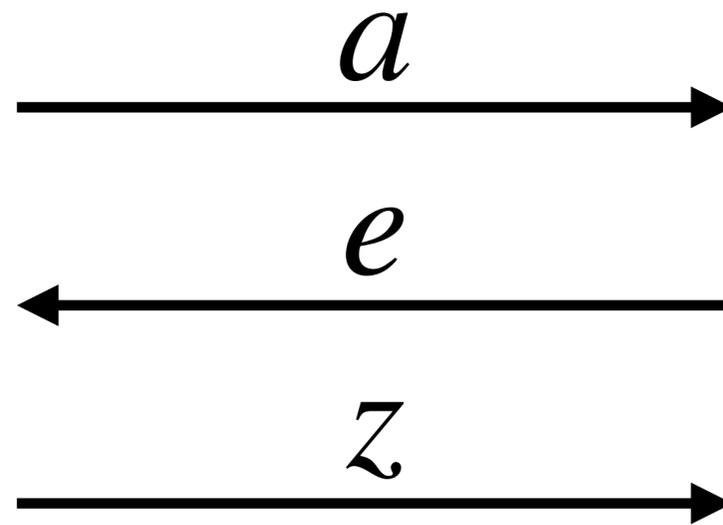
$e_1 \neq e_2$ OR $z_1 \neq z_2$

Tightening Conditions for Extraction

[This work]


 $P(X, w)$


 $V(X)$



$\text{Verify}(a, e, z)$

Strong 2-special soundness:

$w \leftarrow \text{Ext}(X, a, (e_1, z_1), (e_2, z_2))$ such that $R(X, w) = 1$

$e_1 \neq e_2$ OR $z_1 \neq z_2$

...are we done?

What about Zero-knowledge?

- Interestingly, Fischlin's proof of Zero-knowledge also depends on quasi-unique responses
- Unlike extraction, it is not intuitive as to why (or whether it's even necessary)
- [**This work**]: In the absence of unique responses, an explicit attack on *Witness Indistinguishability*

The Attack

- **Fact 1**: In some Sigma protocols, the prover's internal state is exposed to an adversary who has the witness.
eg. Schnorr: $z = xe + r$; given x can solve for r
- **Fact 2**: Once a is fixed, Fischlin's compiler is deterministic

The Attack

- **Fact 1:** In some Sigma protocols, the prover's internal state is exposed to an adversary who has the witness.
eg. Schnorr: $z = xe + r$; given x can solve for r

- **Fact 2:** Once a is fixed, Fischlin's compiler is deterministic

Attacker \mathcal{A} : Given Fischlin-compiled proof π , retrieve prover's internal state, and retrace its steps, i.e. attempt to recompute the proof

The Attack

- **Fact 1:** In some Sigma protocols, the prover's internal state is exposed to an adversary who has the witness.
eg. Schnorr: $z = xe + r$; given x can solve for r

- **Fact 2:** Once a is fixed, Fischlin's compiler is deterministic

Attacker \mathcal{A} : Given Fischlin-compiled proof π , retrieve prover's internal state, and retrace its steps, i.e. attempt to recompute the proof

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

The Attack

- **Fact 1:** In some Sigma protocols, the prover's internal state is exposed to an adversary who has the witness.
eg. Schnorr: $z = xe + r$; given x can solve for r

- **Fact 2:** Once a is fixed, Fischlin's compiler is deterministic

Attacker \mathcal{A} : Given Fischlin-compiled proof π , retrieve prover's internal state, and retrace its steps, i.e. attempt to recompute the proof

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

If the "wrong" witness is used, w.h.p. \mathcal{A} will output a *different* proof $\pi' \neq \pi$

How to Fix it?

- Can't do anything about Fact 1 and Fact 3, i.e. properties of many natural Sigma protocols
- We can fix Fact 2—Fischlin's compiler can be randomized
- Instead of incrementally stepping through challenges, the Prover can try *random* challenges until an accepting transcript is found
- Retrieving Sigma protocol randomness (via Fact 1) is now insufficient to retrace the Prover's steps

This Work

- We explore two dimensions of Fischlin's NIZKPoK compiler:



Applicability:

Only proven for Sigma protocols with 'quasi-unique responses'
(doesn't include logical OR, Pedersen commitment PoK, etc.)

Folklore: "works anyway"

- 1a) Contrary to folklore: attack on Witness Indistinguishability
- 1b) Simple randomization fixes the problem

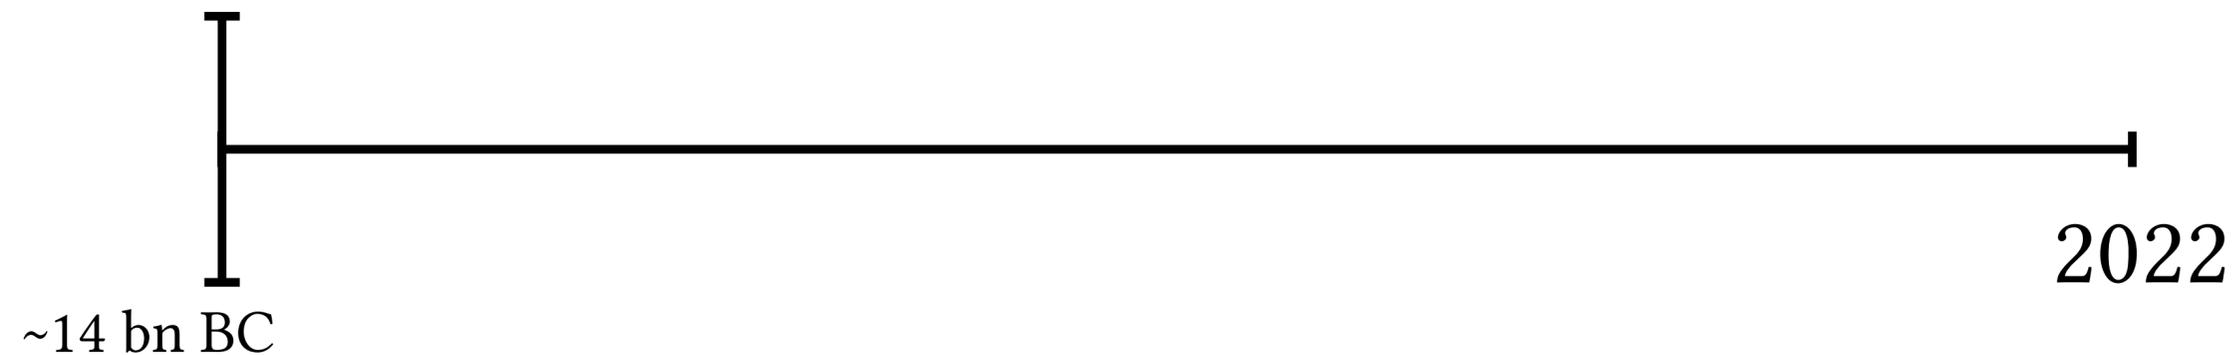
- Computation cost:

Usually the bottleneck — can we improve on it?

- 2) Lower bound: Fischlin05 is optimal up to a small constant
- 3) Application-specific optimization: 200× for EdDSA aggregation

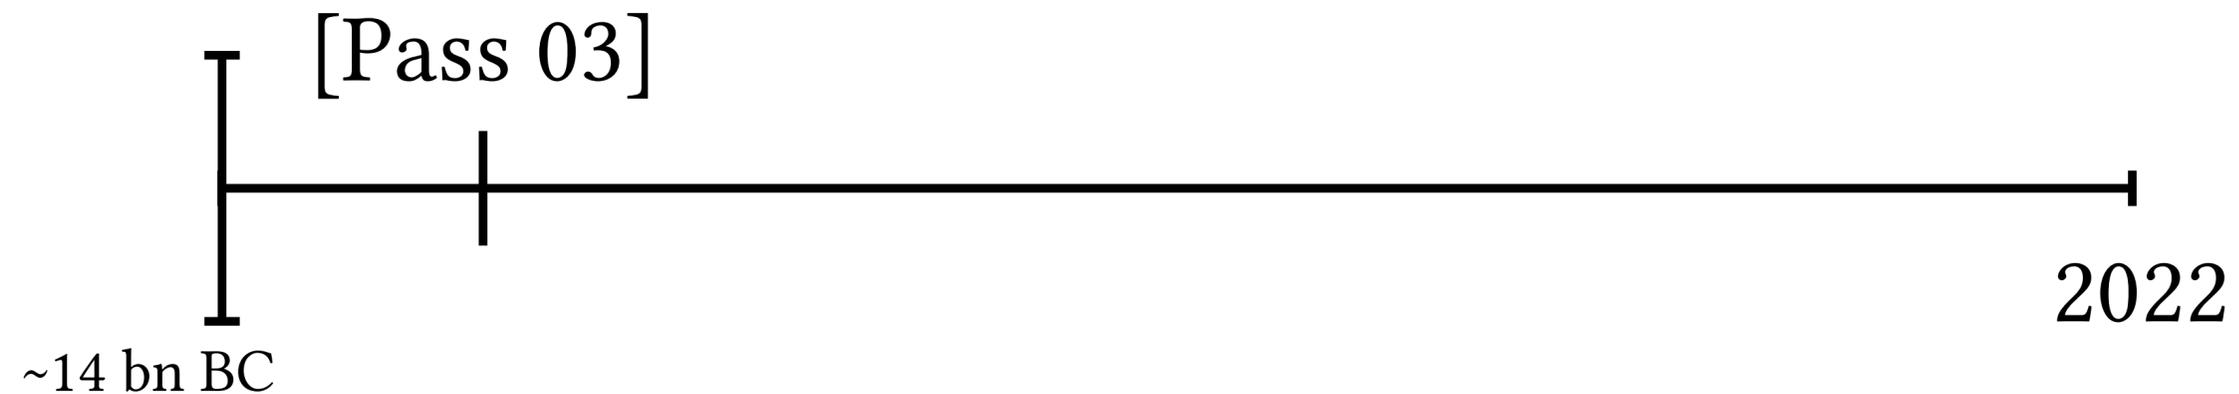
Straight-line Extractable NIZK in the ROM

For simple algebraic statements,
eg. Schnorr's PoK of DLog



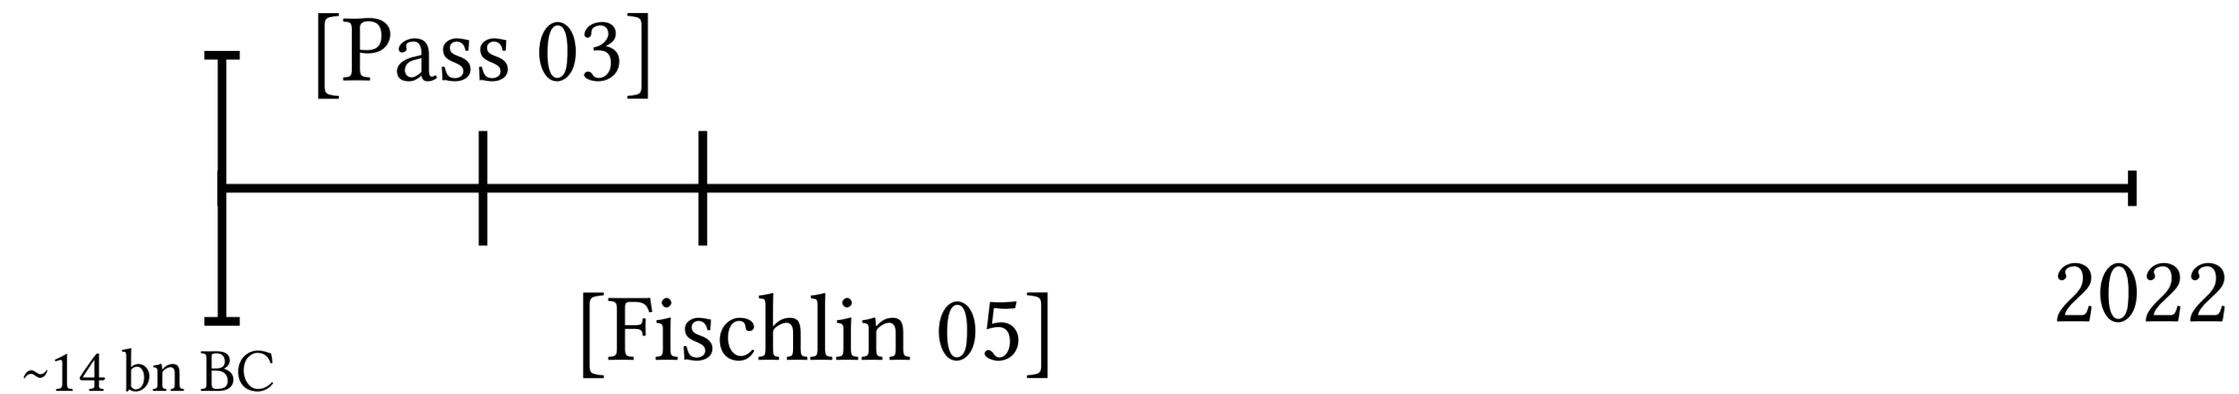
Straight-line Extractable NIZK in the ROM

For simple algebraic statements,
eg. Schnorr's PoK of DLog



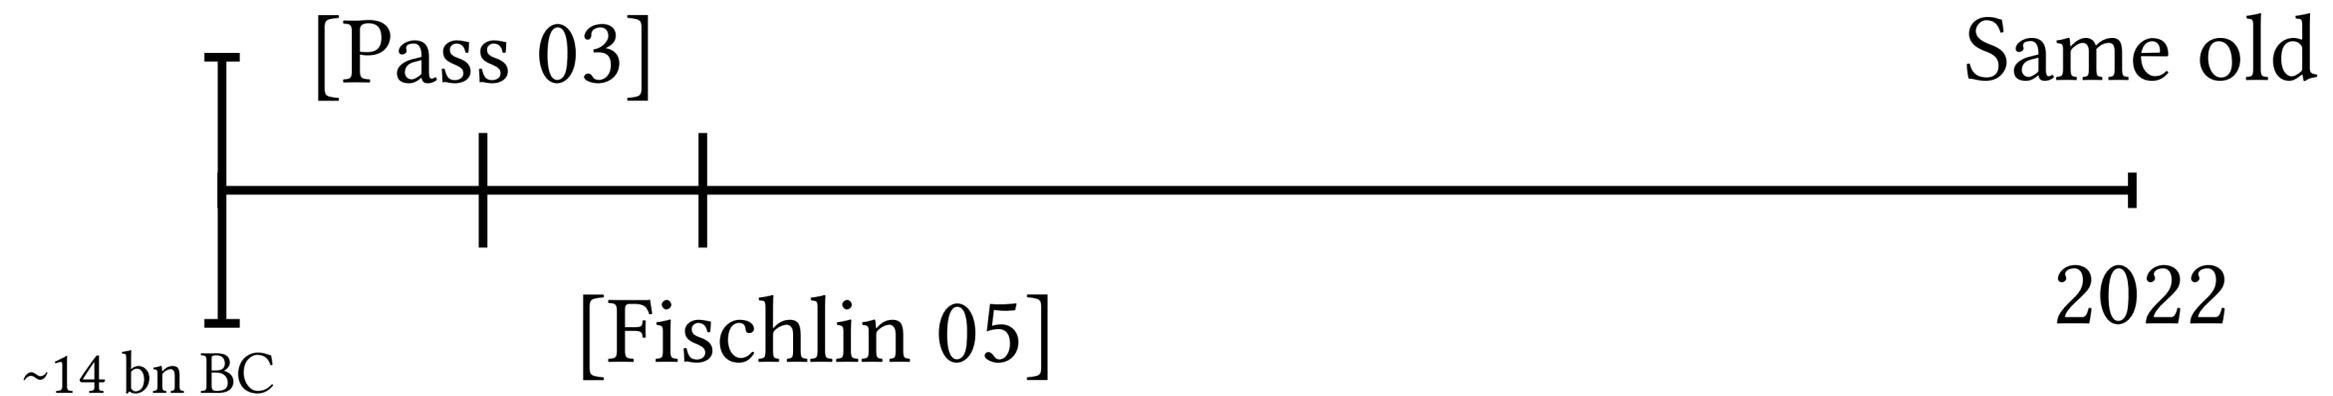
Straight-line Extractable NIZK in the ROM

For simple algebraic statements,
eg. Schnorr's PoK of DLog



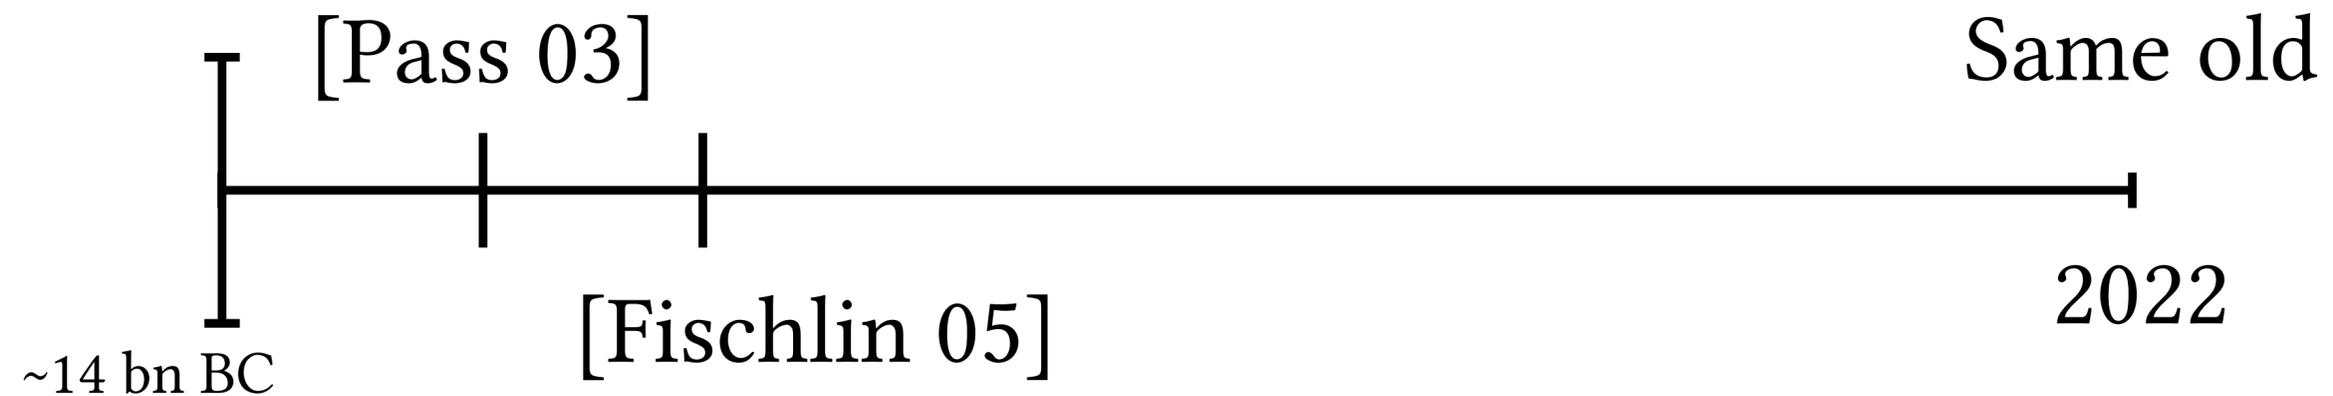
Straight-line Extractable NIZK in the ROM

For simple algebraic statements,
eg. Schnorr's PoK of DLog



Straight-line Extractable NIZK in the ROM

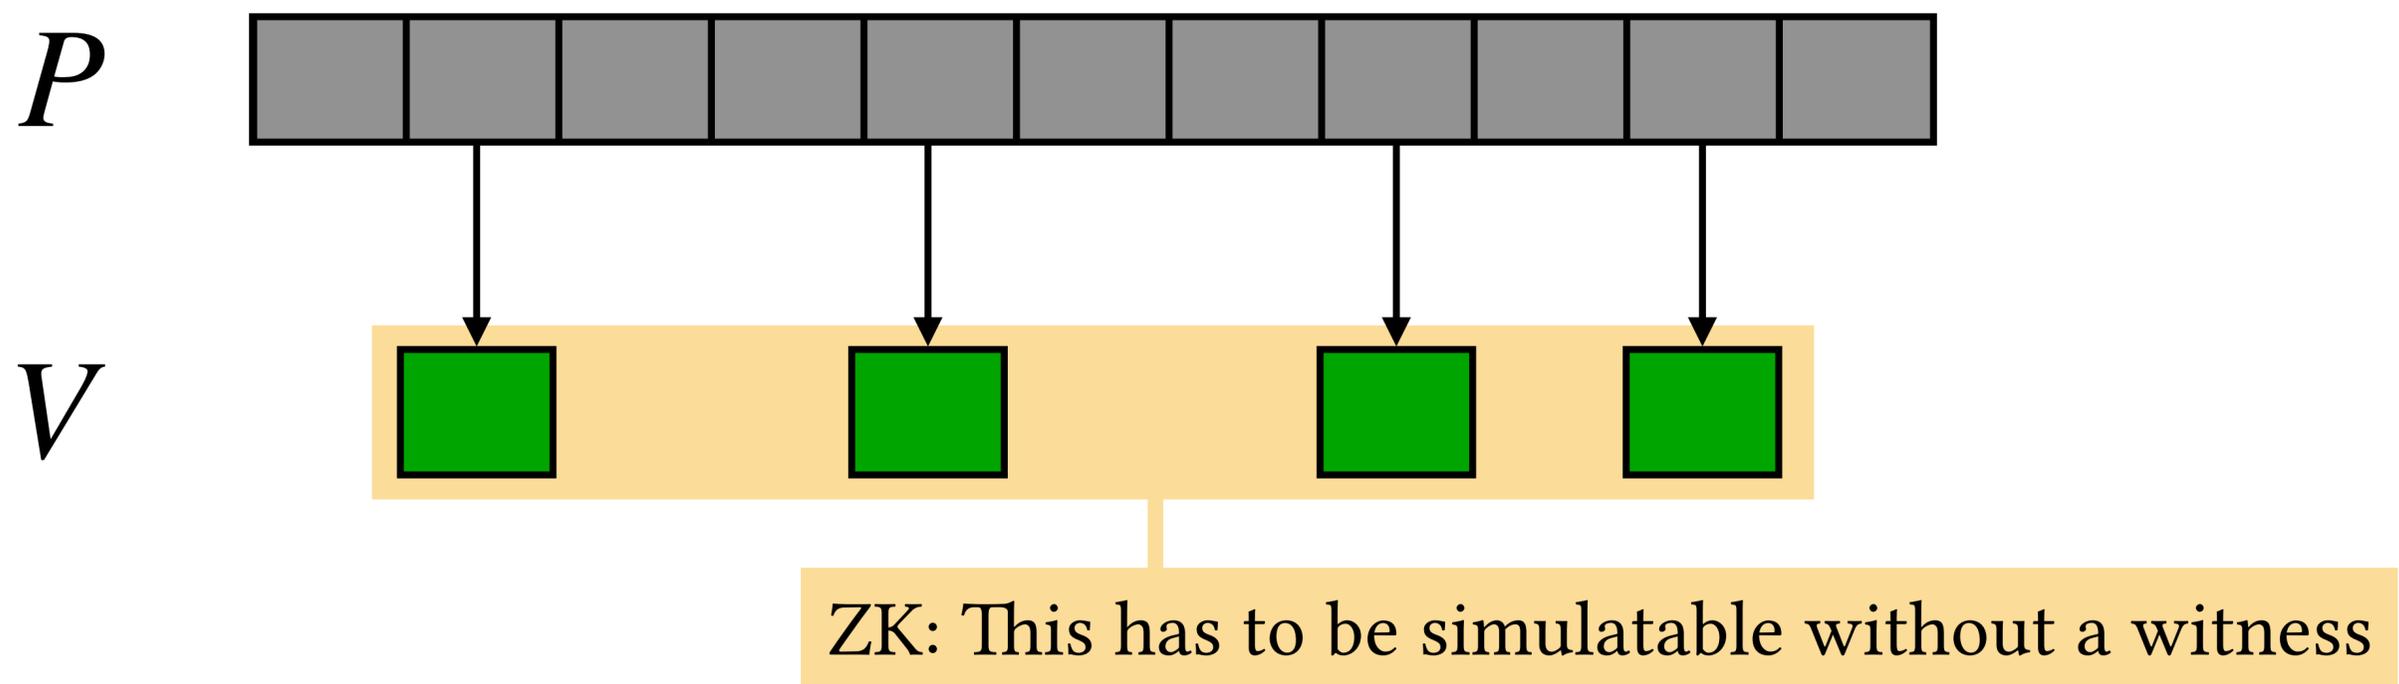
For simple algebraic statements,
eg. Schnorr's PoK of DLog



But why?

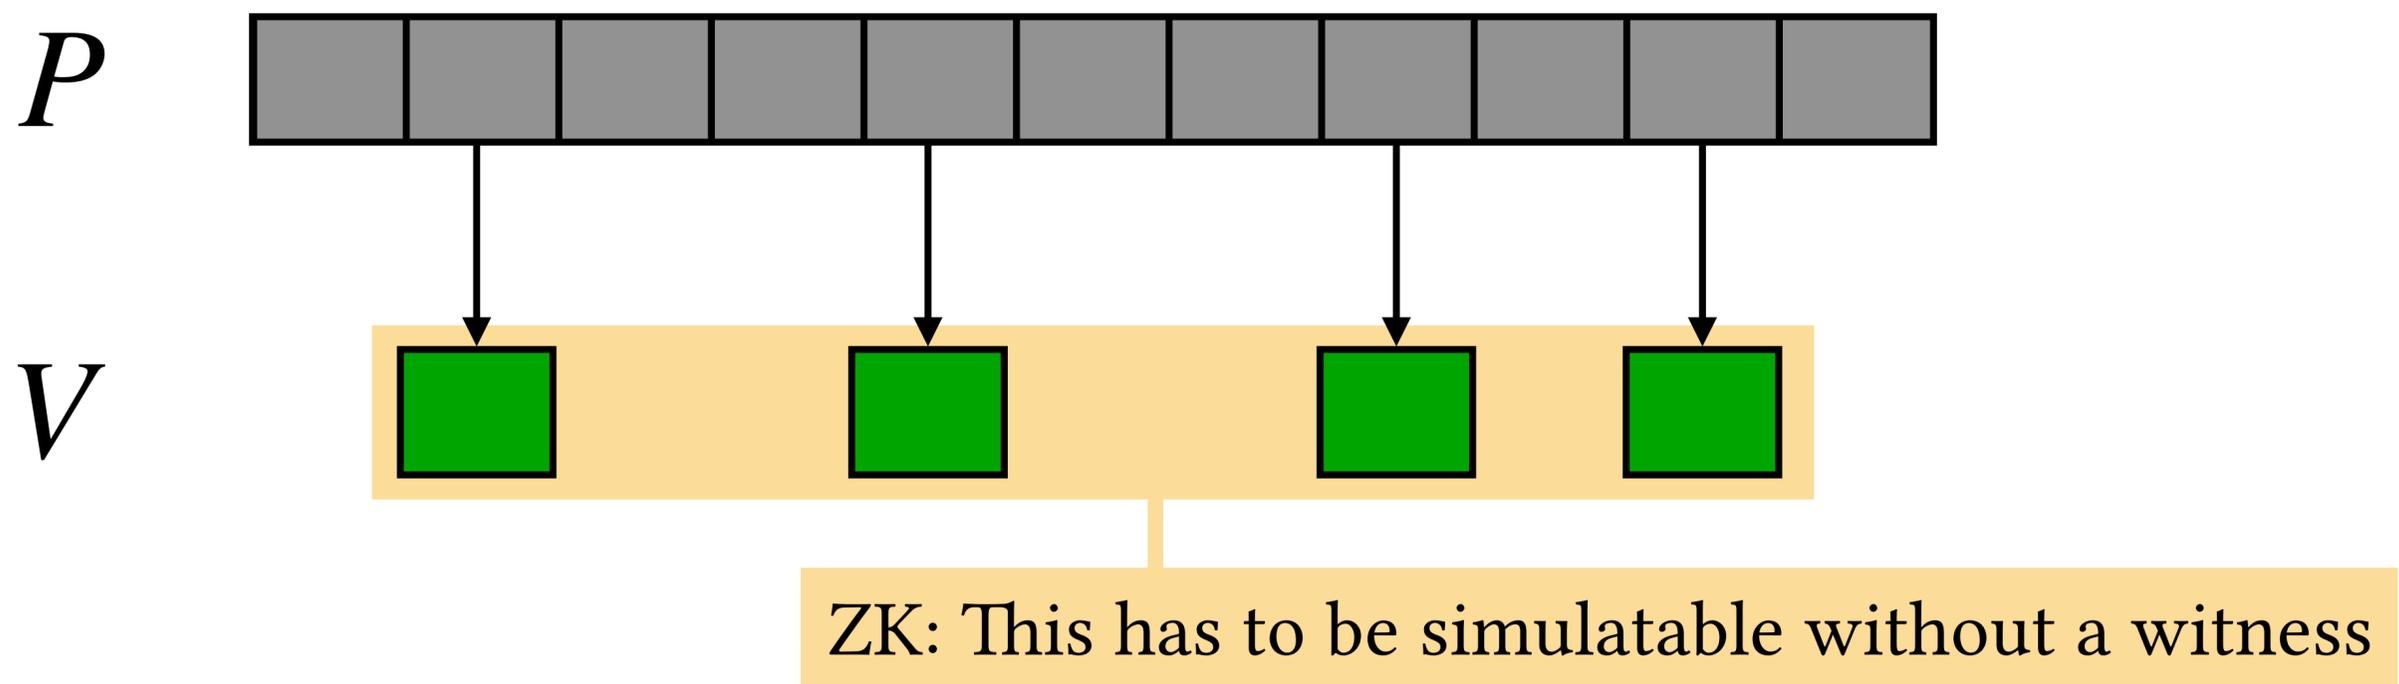
RO Query Complexity for NIZK

- If ZK is desired, we can prove that Fischlin's technique is nearly optimal (within factor of $e \approx 2.7$) for a *non-programming* straight-line extractor
- Our proof is a tightening of an asymptotic bound in [Fischlin 05]
- Lower bound states that if verifier makes V queries and prover P , then $\binom{P}{V} > 2^\kappa$



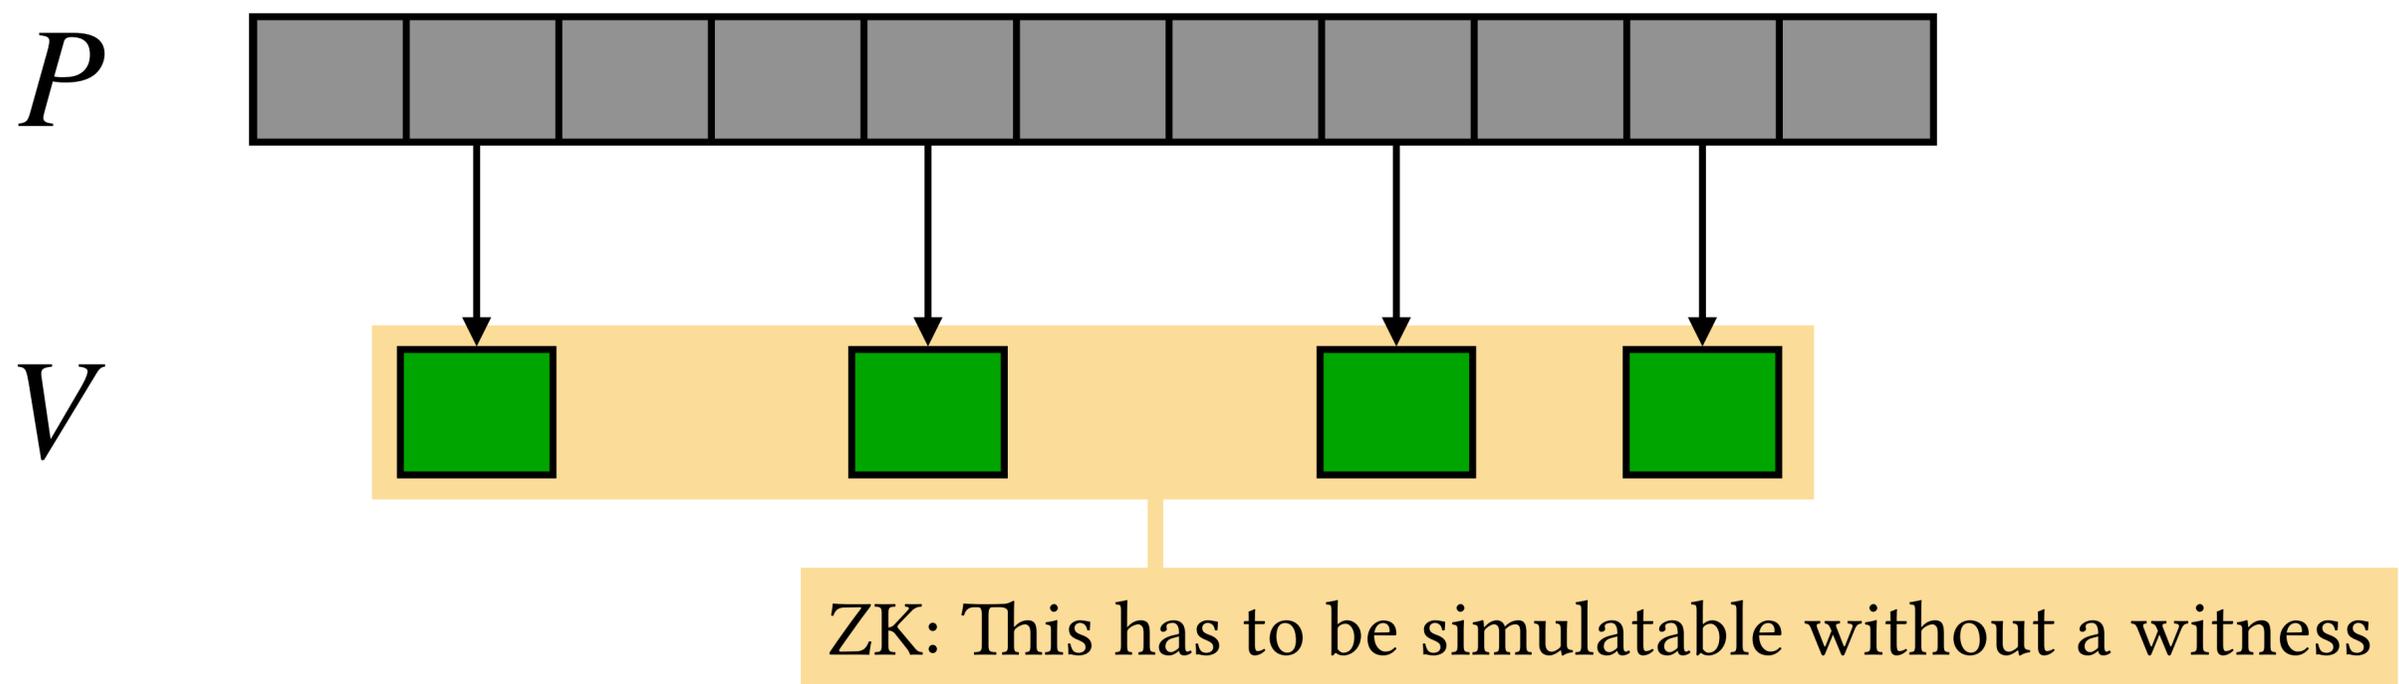
RO Query Complexity for NIZK

- If ZK is desired, we can prove that Fischlin's technique is **nearly** optimal (within factor of $e \approx 2.7$) for a *non-programming* straight-line extractor
- Our proof is a tightening of an asymptotic bound in [Fischlin 05]
- Lower bound states that if verifier makes V queries and prover P , then $\binom{P}{V} > 2^\kappa$



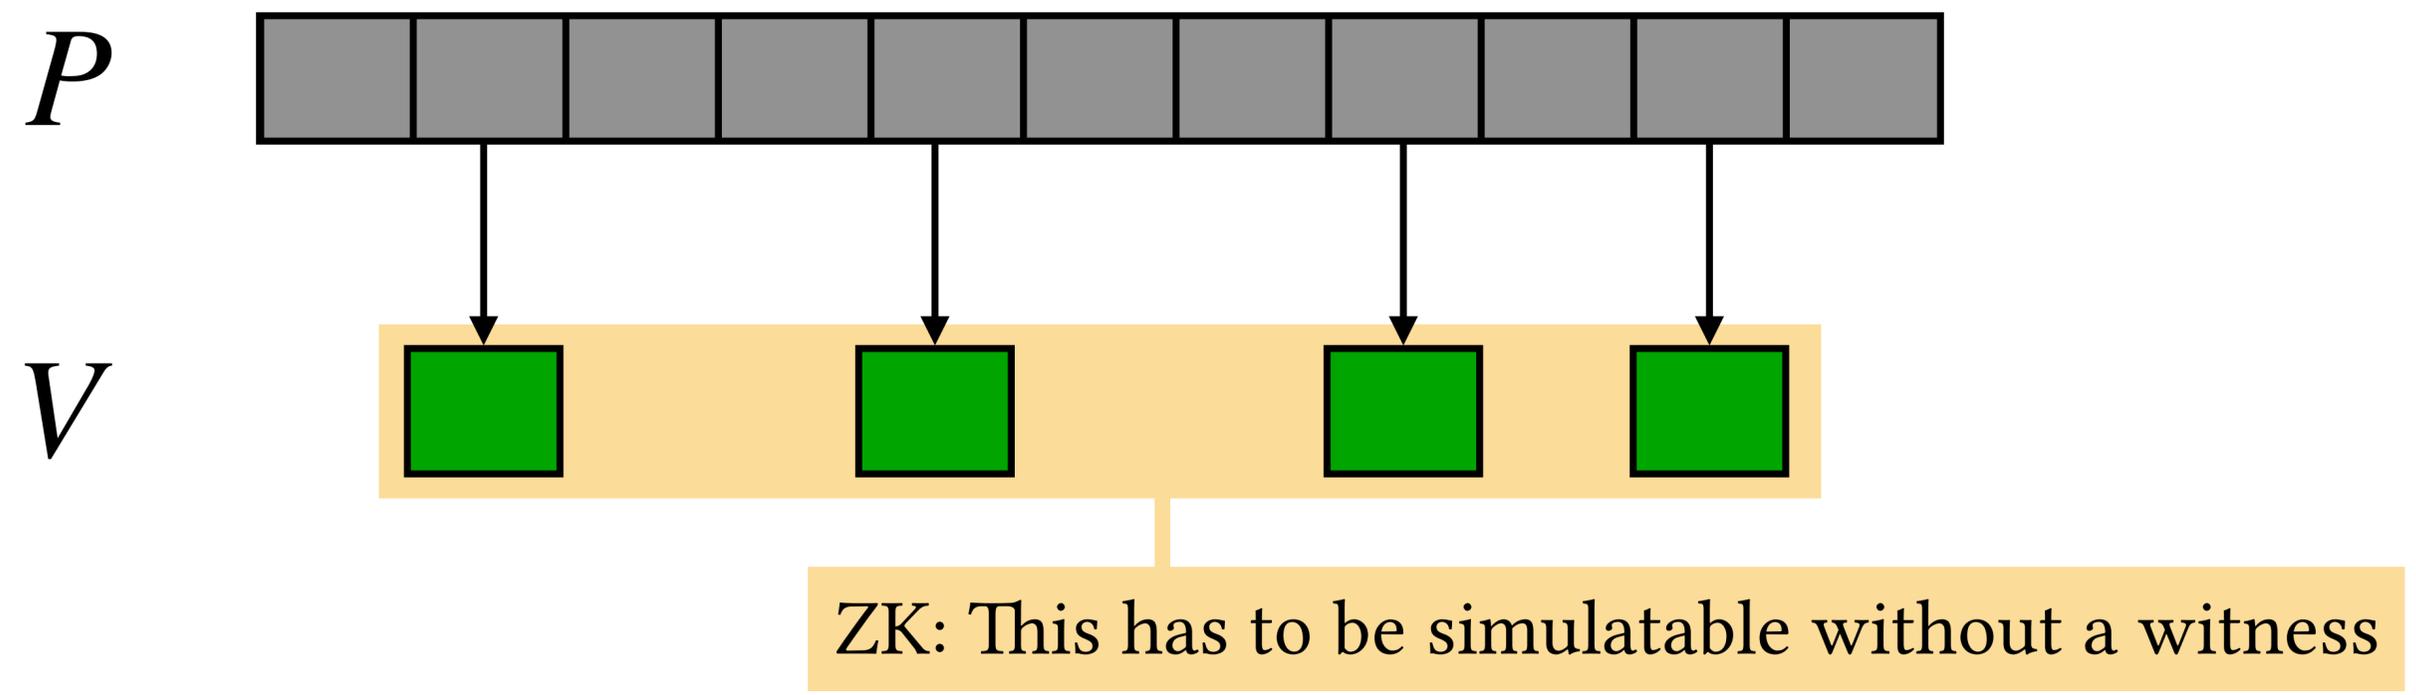
RO Query Complexity for NIZK

- If ZK is desired, we can prove that Fischlin's technique is **nearly** optimal (within factor of $e \approx 2.7$) for a *non-programming* straight-line **Loose bound?** **Or room for improvement?**
- Our proof is a tightening of an asymptotic bound in [Fischlin 05]
- Lower bound states that if verifier makes V queries and prover P , then $\binom{P}{V} > 2^\kappa$



RO Query Complexity for NIZK

- If ZK is desired, we can prove that Fischlin's technique is **nearly** optimal (within factor of $e \approx 2.7$) for a *non-programming* straight-line **Loose bound?** **Or room for improvement?**
- Our proof is a tightening of an asymptotic bound **Having the Prover find *collisions* rather than inversions of H gives a bit of a speedup**
- Lower bound states that if verifier makes V queries

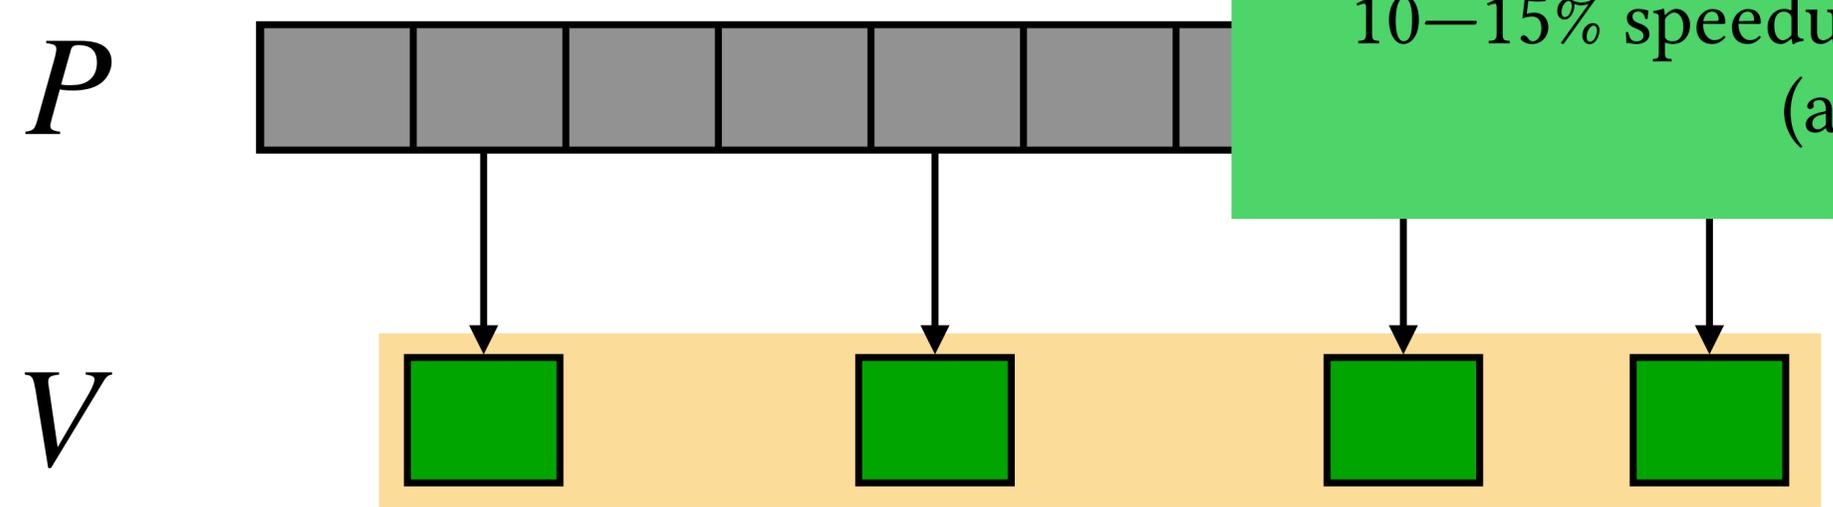


RO Query Complexity for NIZK

- If ZK is desired, we can prove that Fischlin's technique is **nearly** optimal (within factor of $e \approx 2.7$) for a *non-programming* straight-line **Loose bound?** Or room for improvement?
- Our proof is a tightening of an asymptotic bound
- Lower bound states that if verifier makes V queries

Having the Prover find *collisions* rather than inversions of H gives a bit of a speedup

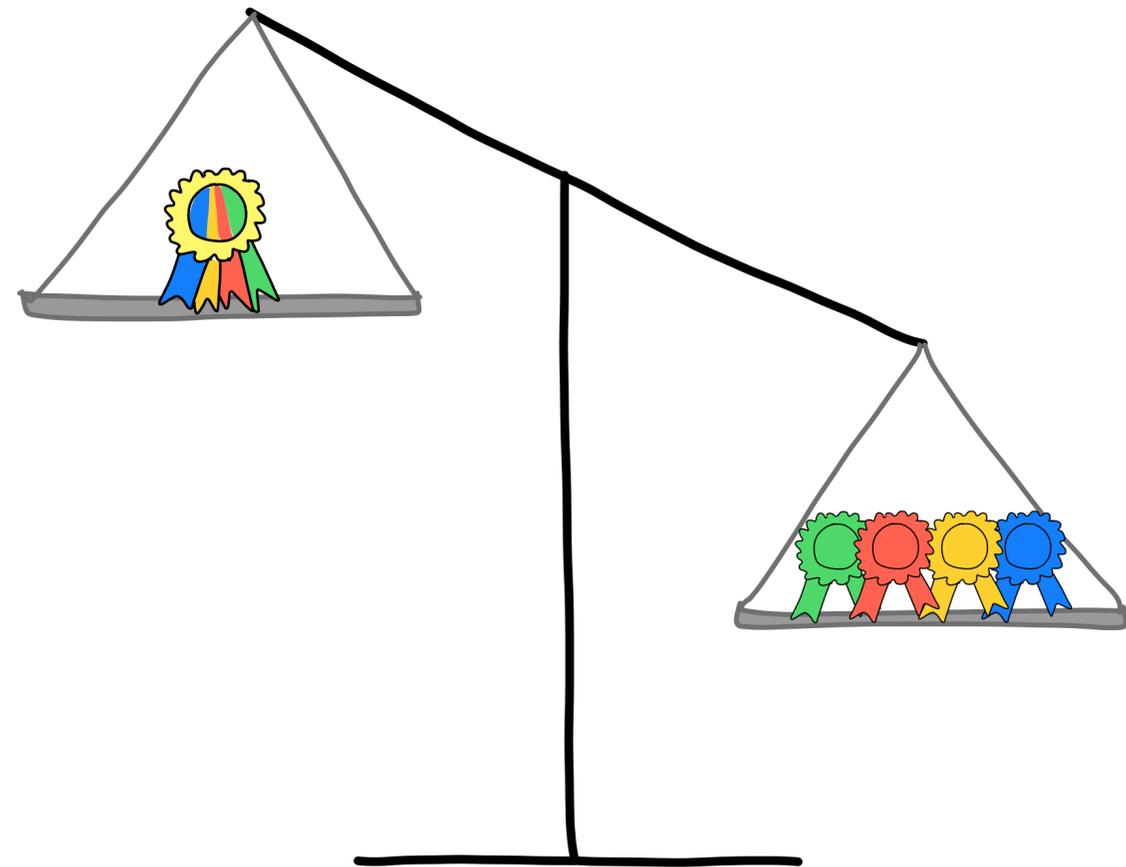
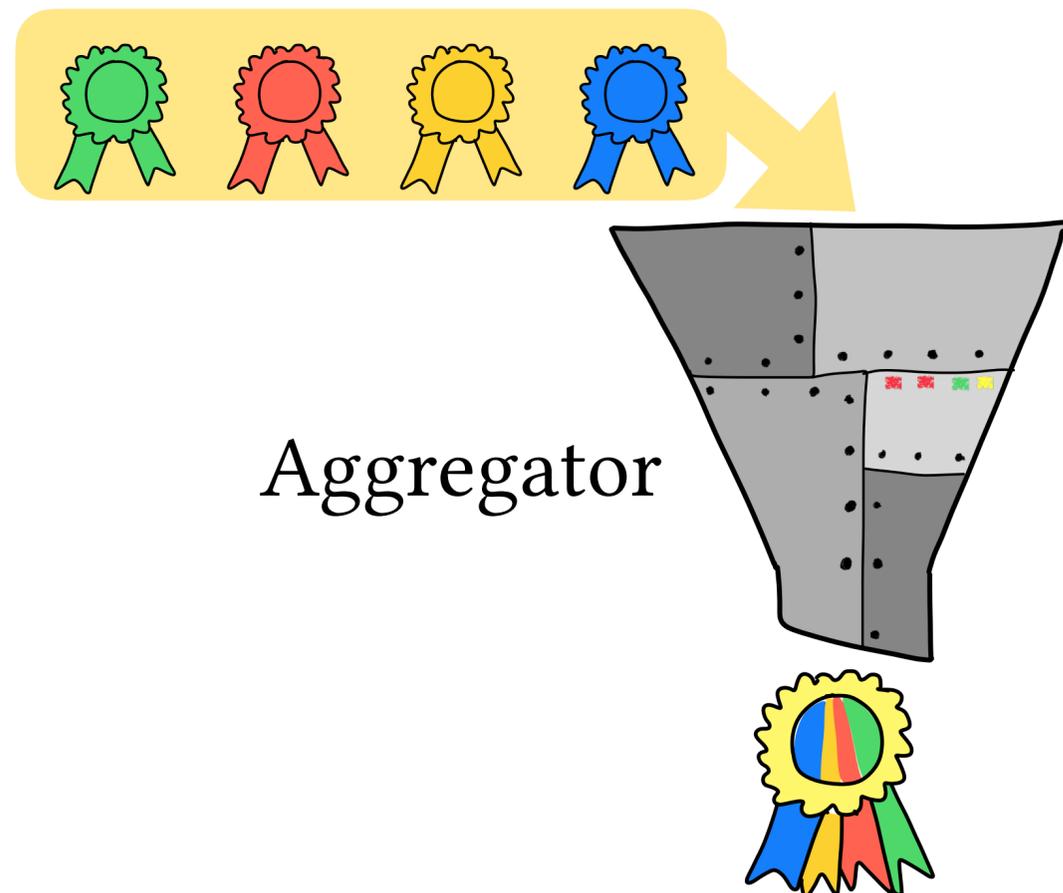
10–15% speedup in the general case for (almost) free



ZK: This has to be simulatable without a witness

Application-Specific Optimization

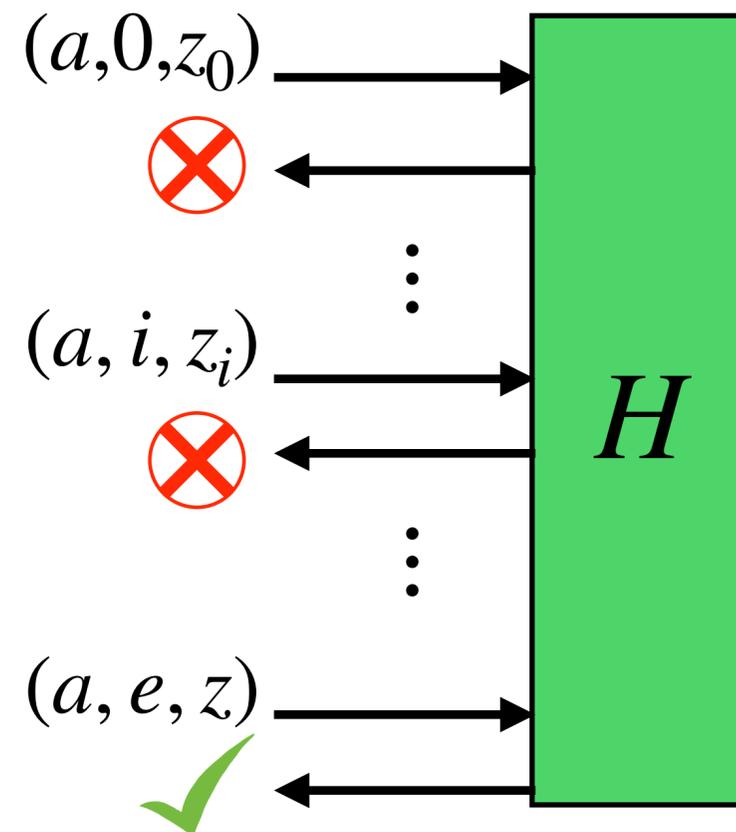
- We show that it is possible to optimize computation cost of Fischlin's technique in specific applications
- We consider Schnorr/EdDSA signature aggregation [CGKN21]: **200× improvement**



Understanding Computation Cost

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

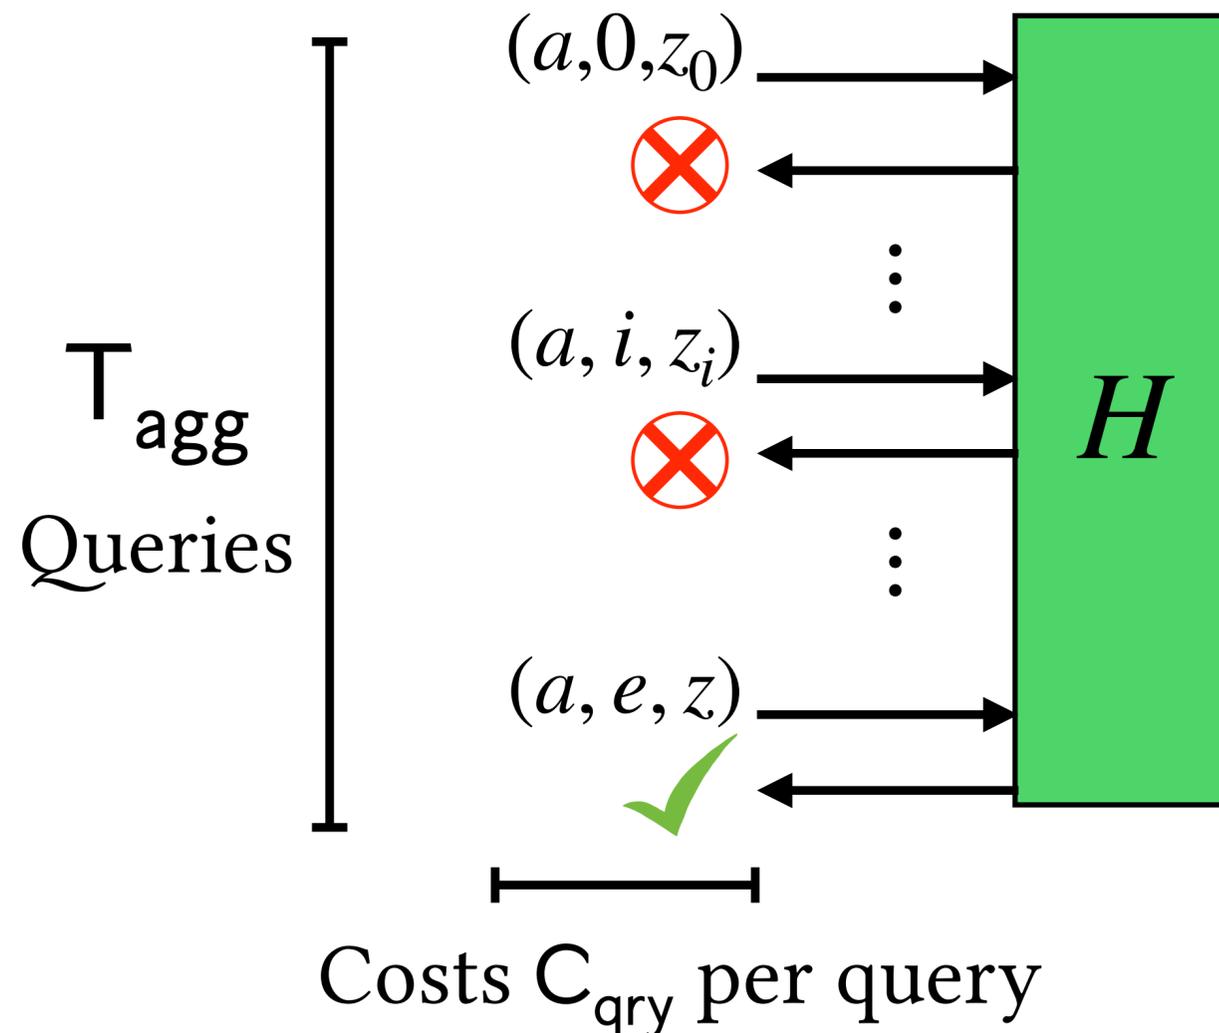
$P(X, w)$: Sample Σ -protocol first message 'a'



Understanding Computation Cost

- Let $H: \{0,1\}^* \mapsto \{0,1\}^\ell$ be a random oracle

$P(X, w)$: Sample Σ -protocol first message 'a'



Total cost: $T_{\text{agg}} \cdot C_{\text{qry}}$

We improve both dimensions

Improving T_{agg}

- The query complexity T_{agg} corresponds to the (expected) running time of finding r inversions of an ℓ -bit hash function
- Insight: finding r **collision** of ℓ' -bit hash is 1.5–2× **faster than inversion**
via birthday attack combinatorial analyses [von Mises 39, Preneel 93]
(ℓ' adjusted to respect the security constraint for the same κ)
- This translates to the Zero-Knowledge (NIZK) setting as well

Improving T_{agg}

- The query complexity T_{agg} corresponds to the (expected) running time of finding r inversions of an ℓ -bit hash function
- Insight: finding r **collision** of ℓ' -bit hash is $1.5\text{--}2\times$ **faster than inversion**
via birthday attack combinatorial analyses [von Mises 39, Preneel 93]
(ℓ' adjusted to respect the security constraint for the same κ)
- This translates to the Zero-Knowledge (NIZK) setting as well

Improving C_{qry}

P

V

$$f \in \mathbb{Z}_q[X]$$



$$\overleftarrow{e \in \mathbb{Z}_q}$$

$$\overrightarrow{f(e)}$$

Improving C_{qry}

C_{qry} in Schnorr aggregation Sigma protocol:

P

V

$$f \in \mathbb{Z}_q[X]$$



$$\leftarrow e \in \mathbb{Z}_q$$

$$\rightarrow f(e)$$

Improving C_{qry}

C_{qry} in Schnorr aggregation Sigma protocol:

P

V

$$f \in \mathbb{Z}_q[X]$$



$e \in \mathbb{Z}_q$

C_{qry} is the cost of computing this

$f(e)$

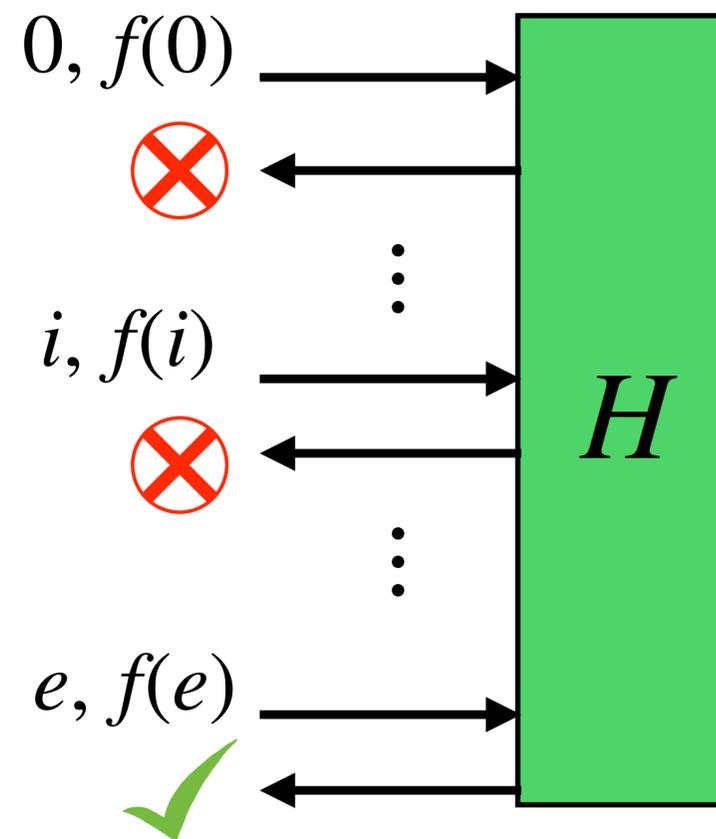
Inherently n multiplications in \mathbb{Z}_q

Improving C_{qry}

C_{qry} in Schnorr aggregation **Fischlin proof**

P

$$f \in \mathbb{Z}_q[X]$$

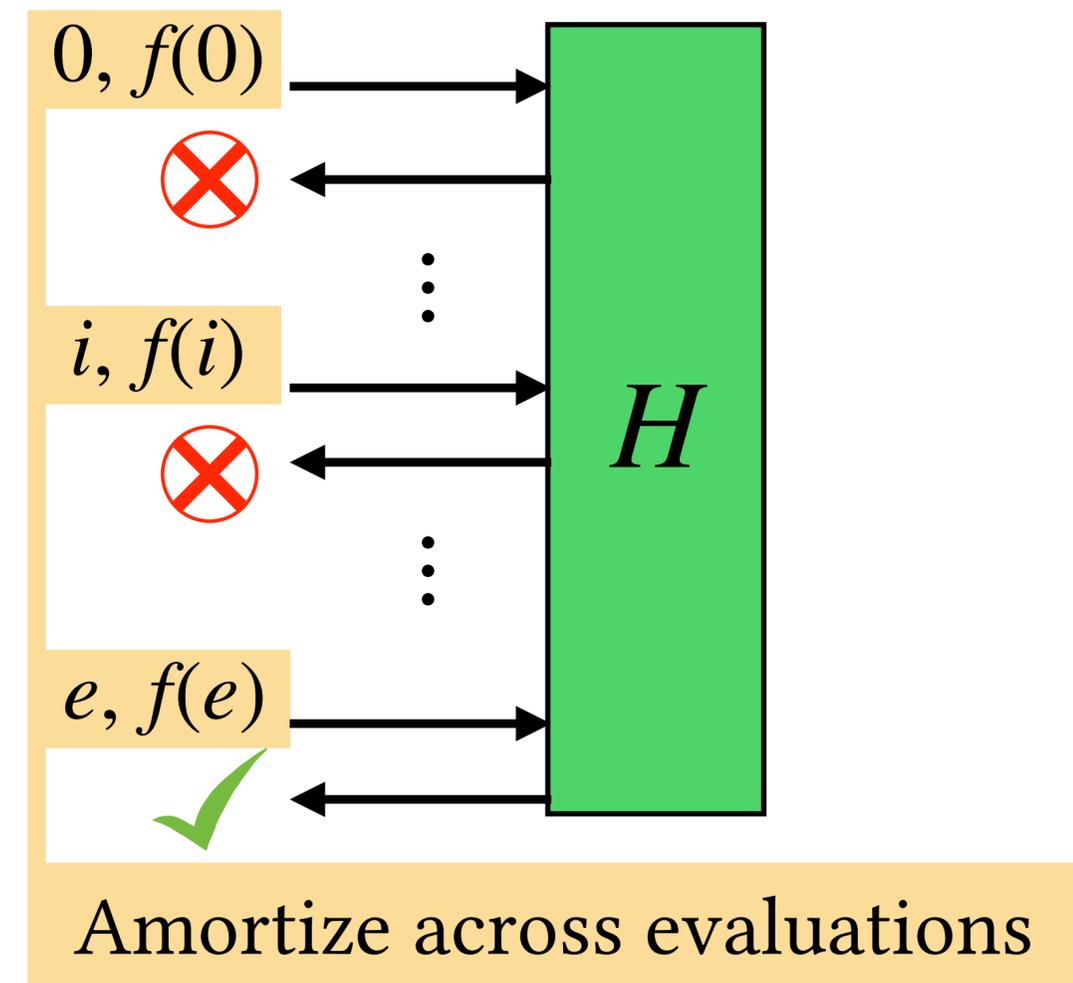


Improving C_{qry}

C_{qry} in Schnorr aggregation **Fischlin proof**

P

$$f \in \mathbb{Z}_q[X]$$



Improving C_{qry}

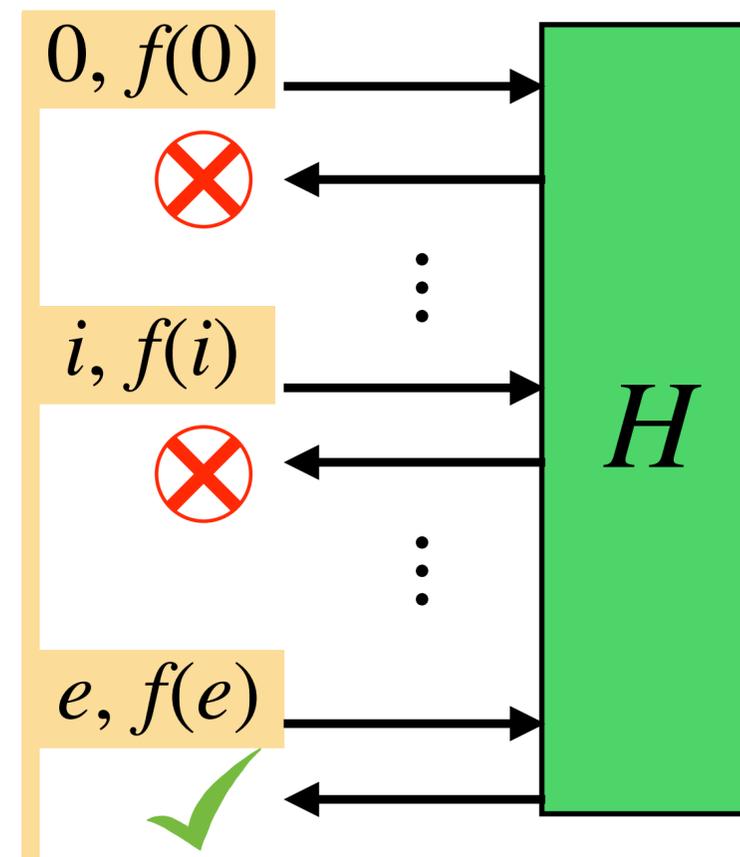
C_{qry} in Schnorr aggregation **Fischlin proof**

P

$$f \in \mathbb{Z}_q[X]$$



FFT, etc.: $O(\log^2(n))$ per eval



Amortize across evaluations

Improving C_{qry}

C_{qry} in Schnorr aggregation **Fischlin proof**

P

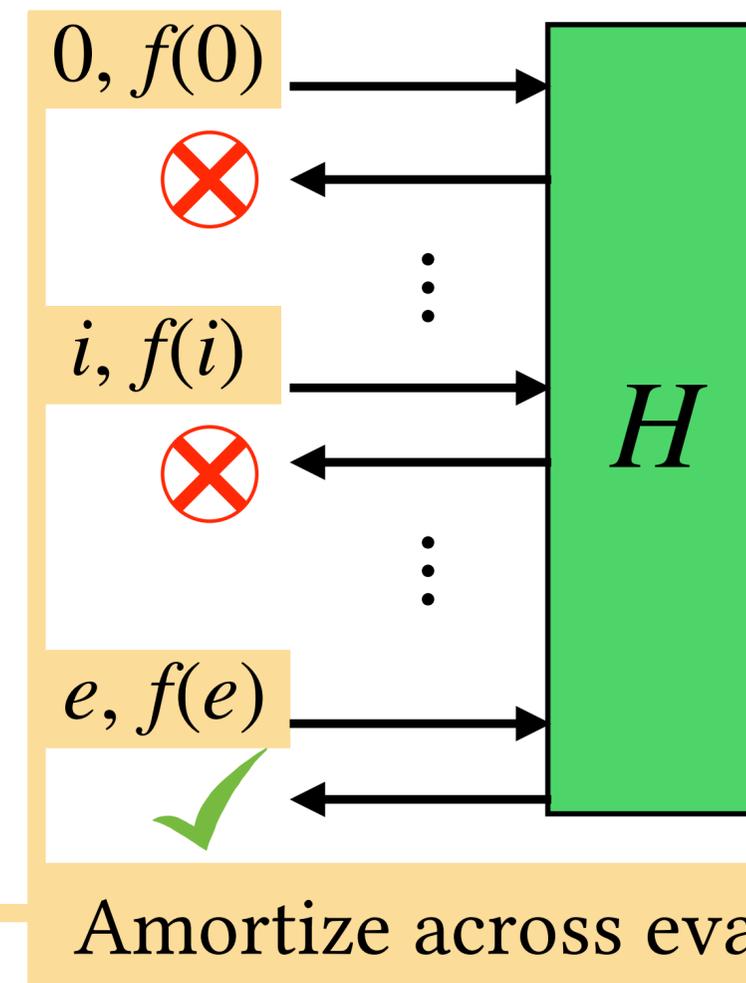
$$f \in \mathbb{Z}_q[X]$$



FFT, etc.: $O(\log^2(n))$ per eval

Most signing curves incompatible with FFT

Asymptotically efficient general multipoint evaluation is unsatisfying for $n < 1000$



Improving C_{qry}

C_{qry} in Schnorr aggregation **Fischlin proof**

P

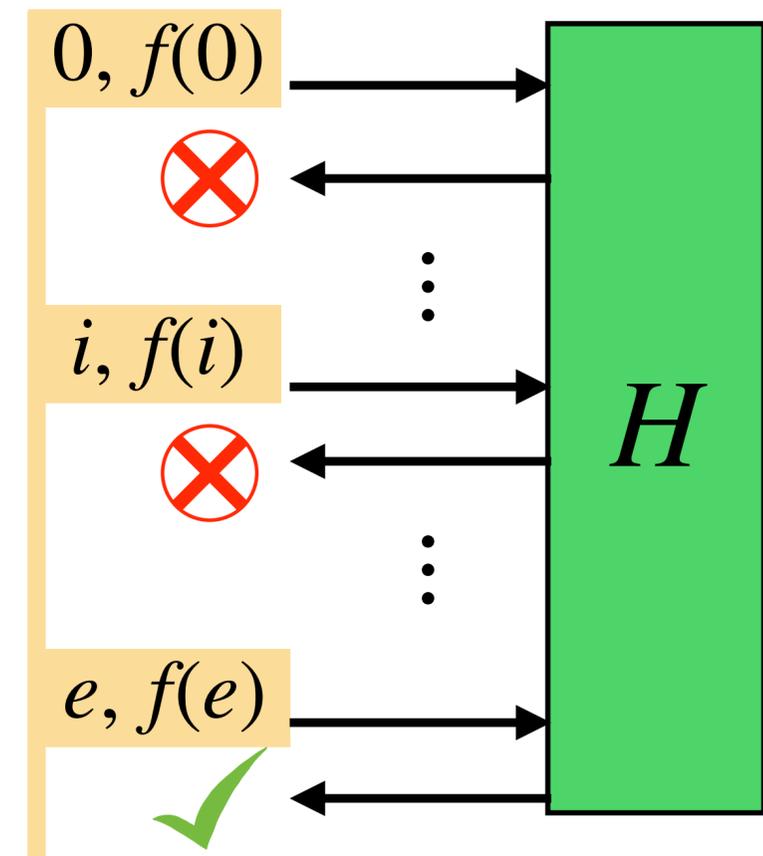
$$f \in \mathbb{Z}_q[X]$$



FFT, etc.: $O(\log^2(n))$ per eval

Most signing curves incompatible with FFT

Asymptotically efficient general multipoint evaluation is unsatisfying for $n < 1000$



Amortize across evaluations

This work: $2\sqrt{n}$ per eval

In Summary

- Fischlin's transform does not preserve Witness Indistinguishability in general — we show how randomization can fix this
- Lower bound explaining lack of progress in SLE in the ROM
 - We show that application-specific optimization is possible
 - Modest general improvement via hash collisions

Thanks!

eprint.iacr.org/2022/393

The Attack

- **Fact 3**: In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

Common a

$P_{\text{OR}}(w_0):$

$P_{\text{OR}}(w_1):$

(e, z)

If $P_{\text{OR}}(w_0)$ and $P_{\text{OR}}(w_1)$ “agree” at e , then they “disagree” at any $e' \neq e$

The Attack

- **Fact 3**: In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

$P_{\text{OR}}(w_0):$

Common a

If $P_{\text{OR}}(w_0)$ and $P_{\text{OR}}(w_1)$ “agree” at e , then they “disagree” at any $e' \neq e$

$P_{\text{OR}}(w_1):$

(e, z)

The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

Common a

$P_{\text{OR}}(w_0):$

$(0, z_0)$

$(1, z_1)$

\vdots

$P_{\text{OR}}(w_1):$

(e, z)

If $P_{\text{OR}}(w_0)$ and $P_{\text{OR}}(w_1)$ “agree” at e , then they “disagree” at any $e' \neq e$

The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

Common a

$P_{\text{OR}}(w_0):$

$(0, z_0)$

$(1, z_1)$

\vdots

$P_{\text{OR}}(w_1): (0, z'_0) (1, z'_1) \dots (e, z)$

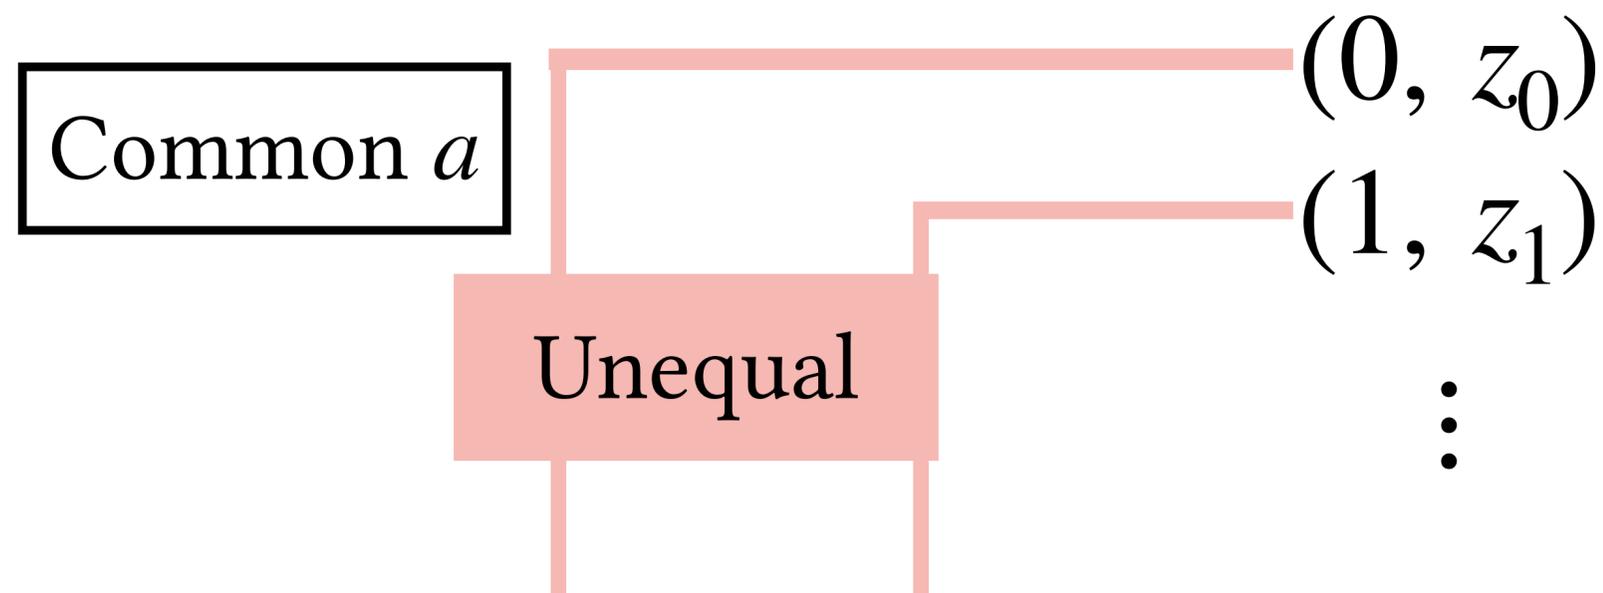
If $P_{\text{OR}}(w_0)$ and $P_{\text{OR}}(w_1)$ “agree” at e , then they “disagree” at any $e' \neq e$

The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

$P_{\text{OR}}(w_0):$



$P_{\text{OR}}(w_1): (0, z'_0) (1, z'_1) \dots (e, z)$

If $P_{\text{OR}}(w_0)$ and $P_{\text{OR}}(w_1)$ “agree” at e , then they “disagree” at any $e' \neq e$

The Attack

- **Fact 3**: In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

$P_{\text{OR}}(w_0):$

$(0, z_0)$

$(1, z_1)$

\vdots

$P_{\text{OR}}(w_1): (0, z'_0) (1, z'_1) \dots (e, z)$

The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

$P_{\text{OR}}(w_0):$

$(0, z_0)$

$(1, z_1)$

\vdots

H

$P_{\text{OR}}(w_1): (0, z'_0) (1, z'_1) \dots (e, z)$

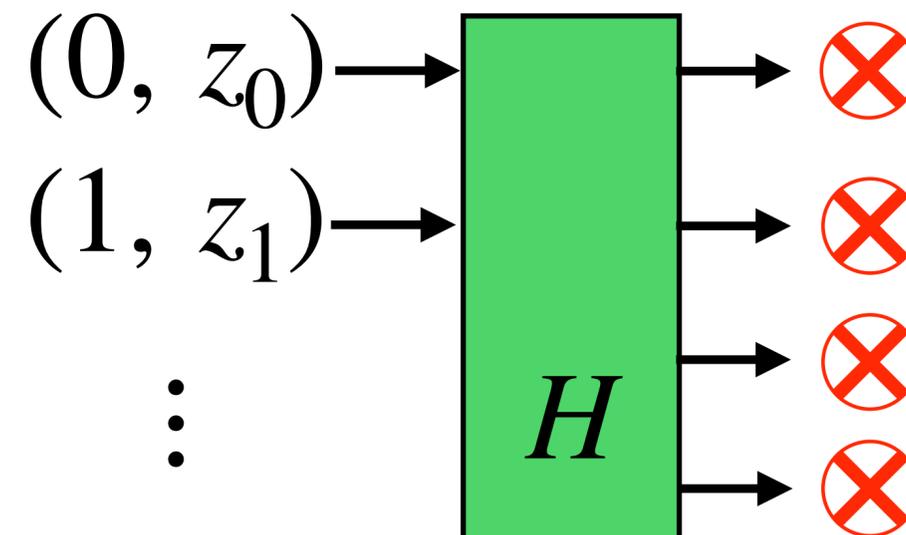
The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

$P_{\text{OR}}(w_0):$



$P_{\text{OR}}(w_1):$ $(0, z'_0)$ $(1, z'_1)$... $(e, z) \rightarrow$ \checkmark

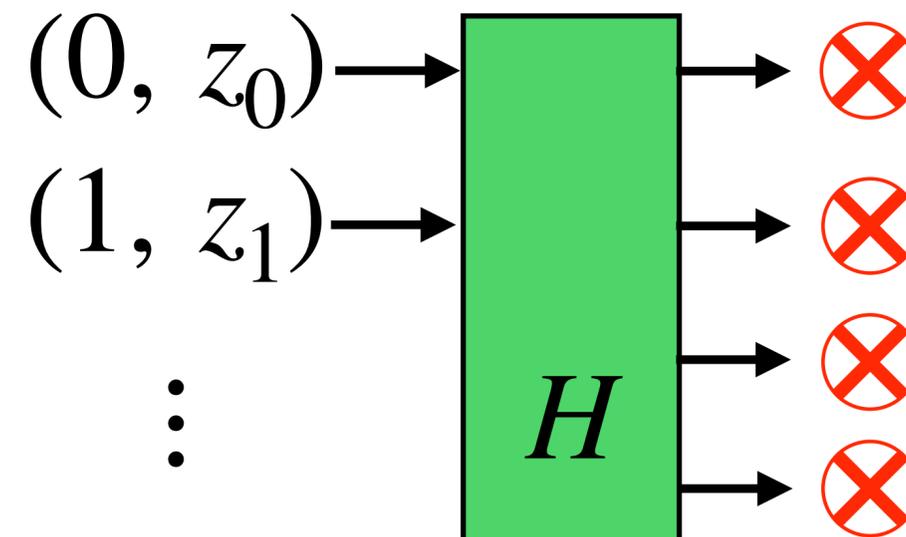
The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

$P_{\text{OR}}(w_0)$:



W.h.p., only one path—
induced by one
of $P_{\text{OR}}(w_0)$ or
 $P_{\text{OR}}(w_1)$ —
is plausible

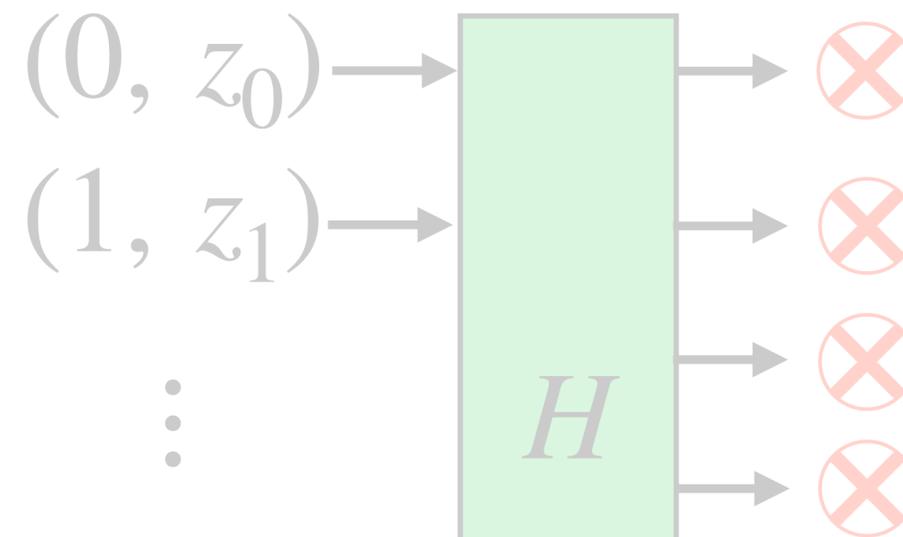
The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

$P_{\text{OR}}(w_0)$:



$P_{\text{OR}}(w_1)$: $(0, z'_0)$ $(1, z'_1)$... $(e, z) \rightarrow$

W.h.p., only one path—
induced by one
of $P_{\text{OR}}(w_0)$ or
 $P_{\text{OR}}(w_1)$ —
is plausible

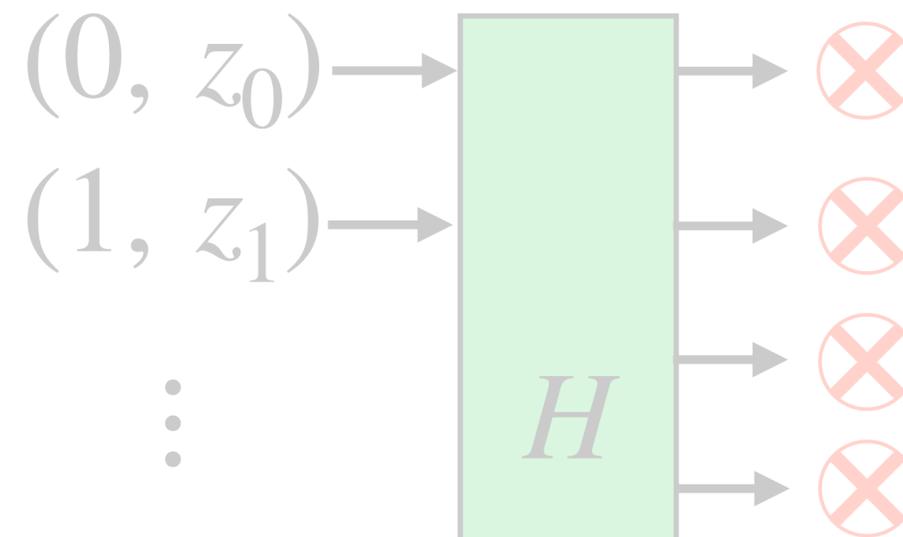
The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

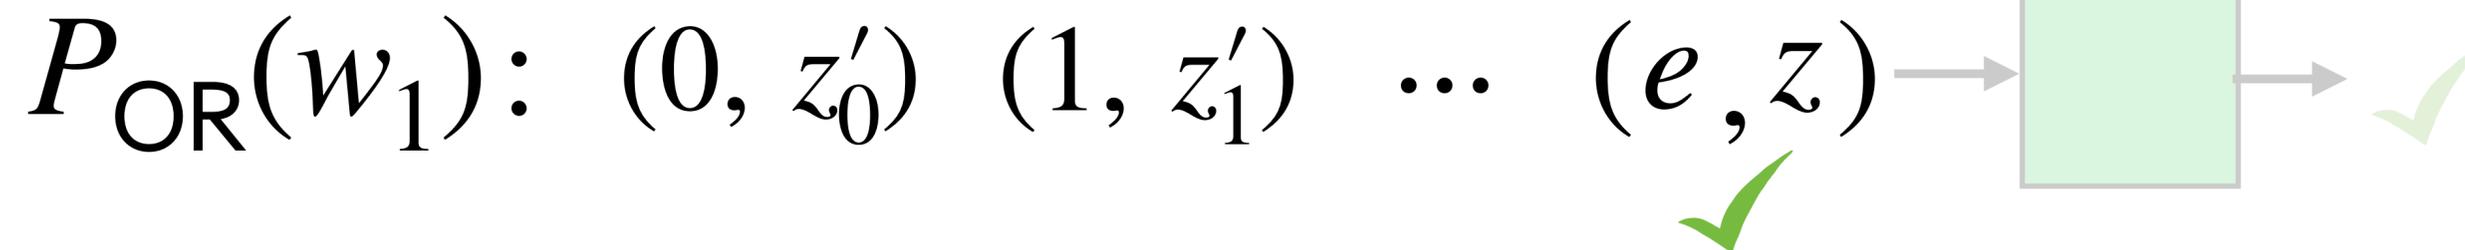
Consider a given (a, e, z)

Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

$P_{\text{OR}}(w_0)$:



W.h.p., only one path—
induced by one
of $P_{\text{OR}}(w_0)$ or
 $P_{\text{OR}}(w_1)$ —
is plausible



The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

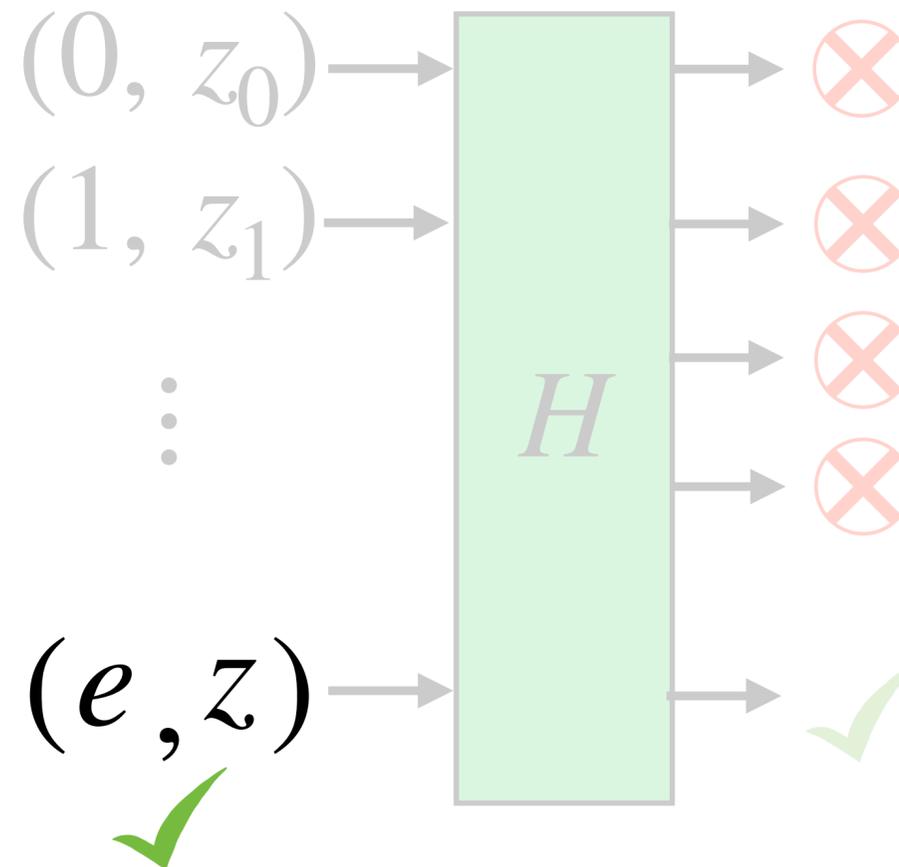
Consider a given (a, e, z)

Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

This path induces fresh queries to H

$P_{\text{OR}}(w_1):$ $(0, z'_0)$ $(1, z'_1)$... (e, z)

$P_{\text{OR}}(w_0):$



W.h.p., only one path—
induced by one
of $P_{\text{OR}}(w_0)$ or
 $P_{\text{OR}}(w_1)$ —
is plausible

The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

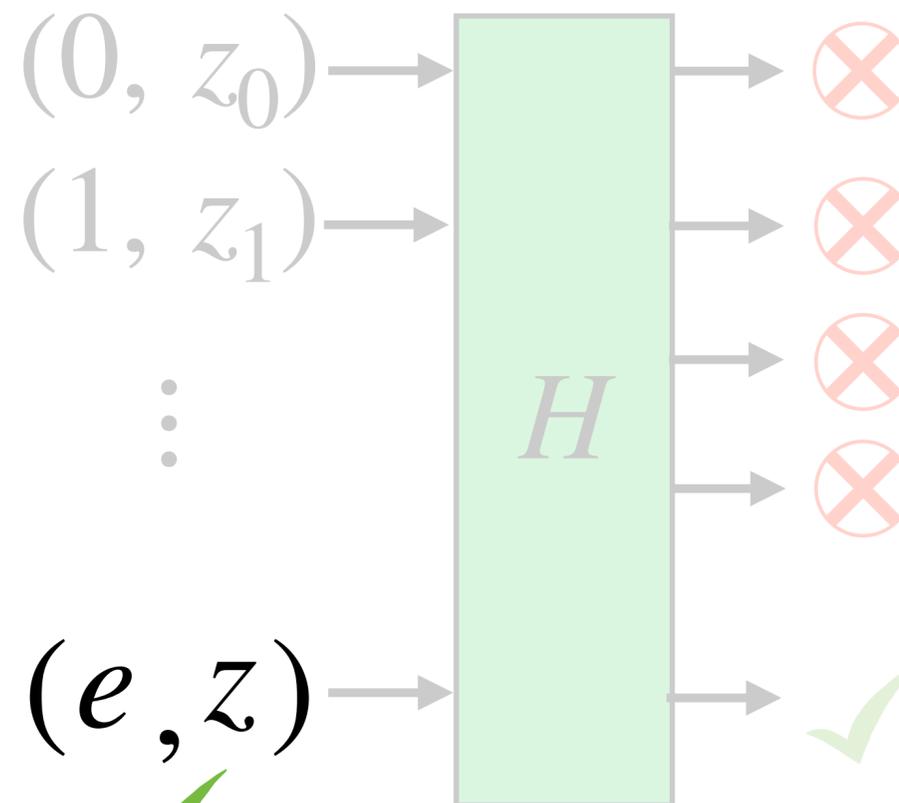
Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

This path induces fresh queries to H

$P_{\text{OR}}(w_1):$ $(0, z'_0)$ $(1, z'_1)$... (e, z)

$P_{\text{OR}}(w_0):$



W.h.p., only one path—
induced by one
of $P_{\text{OR}}(w_0)$ or
 $P_{\text{OR}}(w_1)$ —
is plausible

The Attack

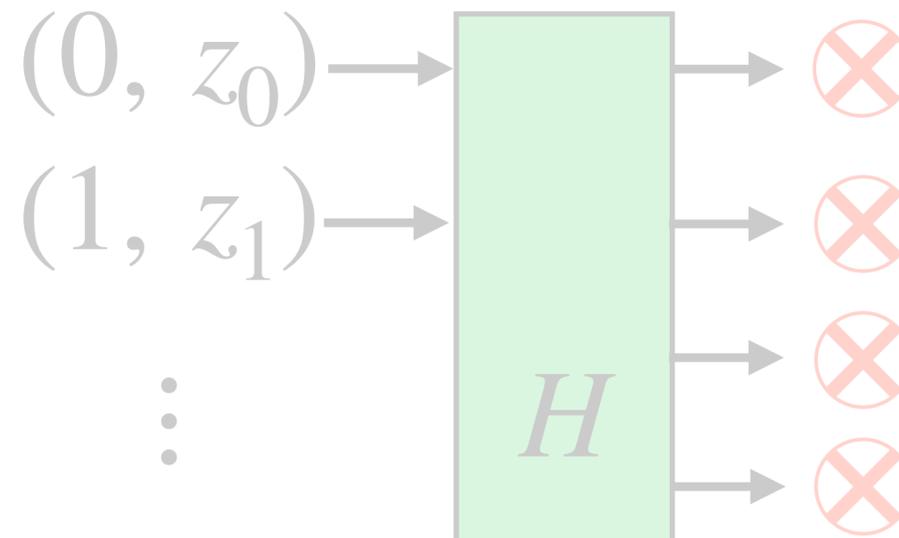
- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

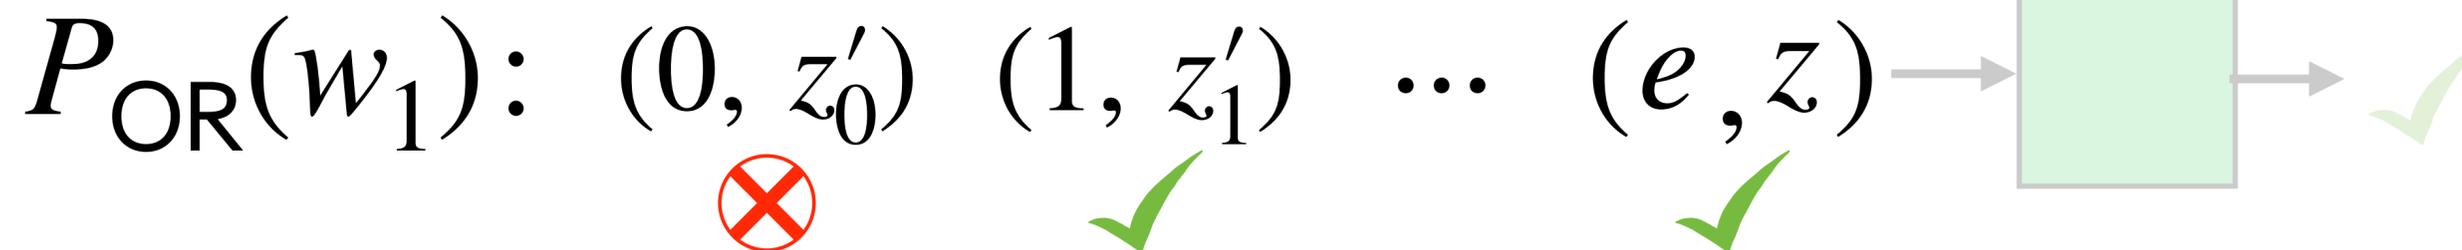
Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

This path induces fresh queries to H

$P_{\text{OR}}(w_0)$:



W.h.p., only one path—
induced by one
of $P_{\text{OR}}(w_0)$ or
 $P_{\text{OR}}(w_1)$ —
is plausible



The Attack

- **Fact 3:** In some Sigma protocols, for the same (a, e) , the response z will depend on which witness is used. e.g. PoK of w_0 OR w_1

Consider a given (a, e, z)

Given (a, e, z) produced by Fischlin's compiler, we can test which path is "plausible"

This path induces fresh queries to H

Would have terminated here

