The Fiat-Shamir Transform

Ron Rothblum

Technion

The Fiat-Shamir Transform

[FS86]

In a nutshell: Awesome technique for minimizing interaction in public-coin interactive protocols.

Fascinating both in theory and in practice.

* Original goal was transforming ID schemes into signature schemes.

Interactive Argument [BCC88]

$x \in L?$

Prover P **Verifier V**

Public-coin if all V does is flip coins and send the result

- **Completeness**: P convinces
- **Computational Soundness** cheating prover can convince V to accept $x \notin L$ (except with negligible probability).

The Fiat-Shamir Transform

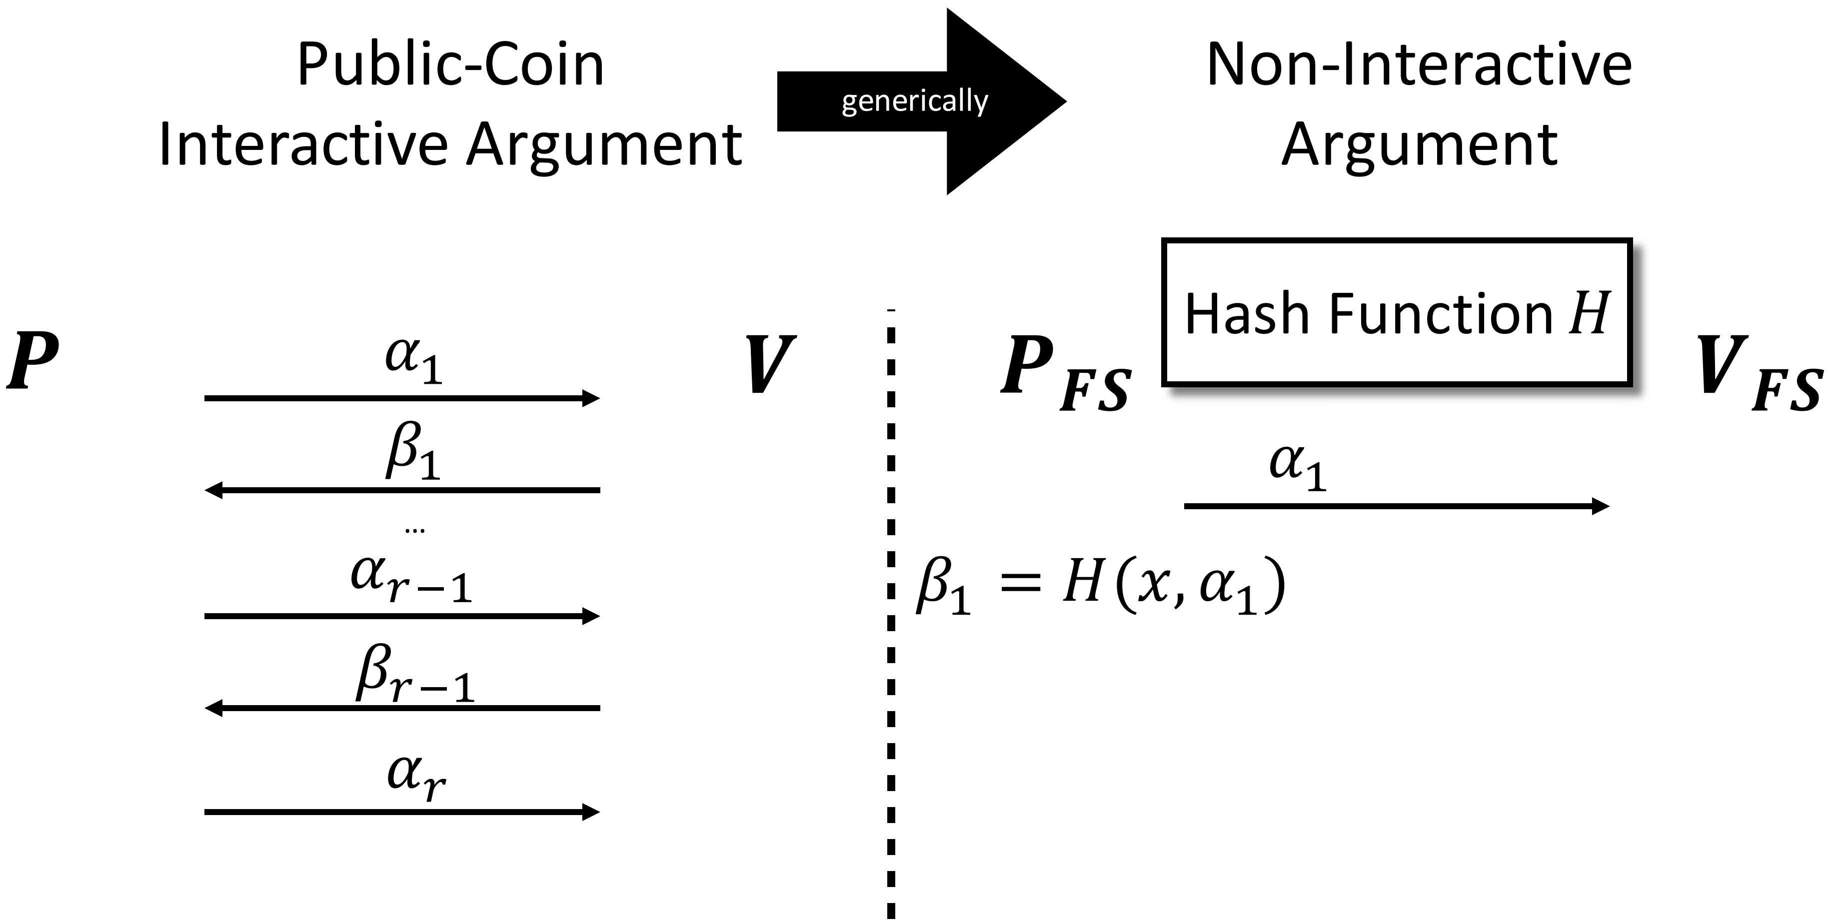

(Each β_i uniformly random)

The Fiat-Shamir Transform

Public-Coin Interactive Argument → (generically) → Non-Interactive Argument

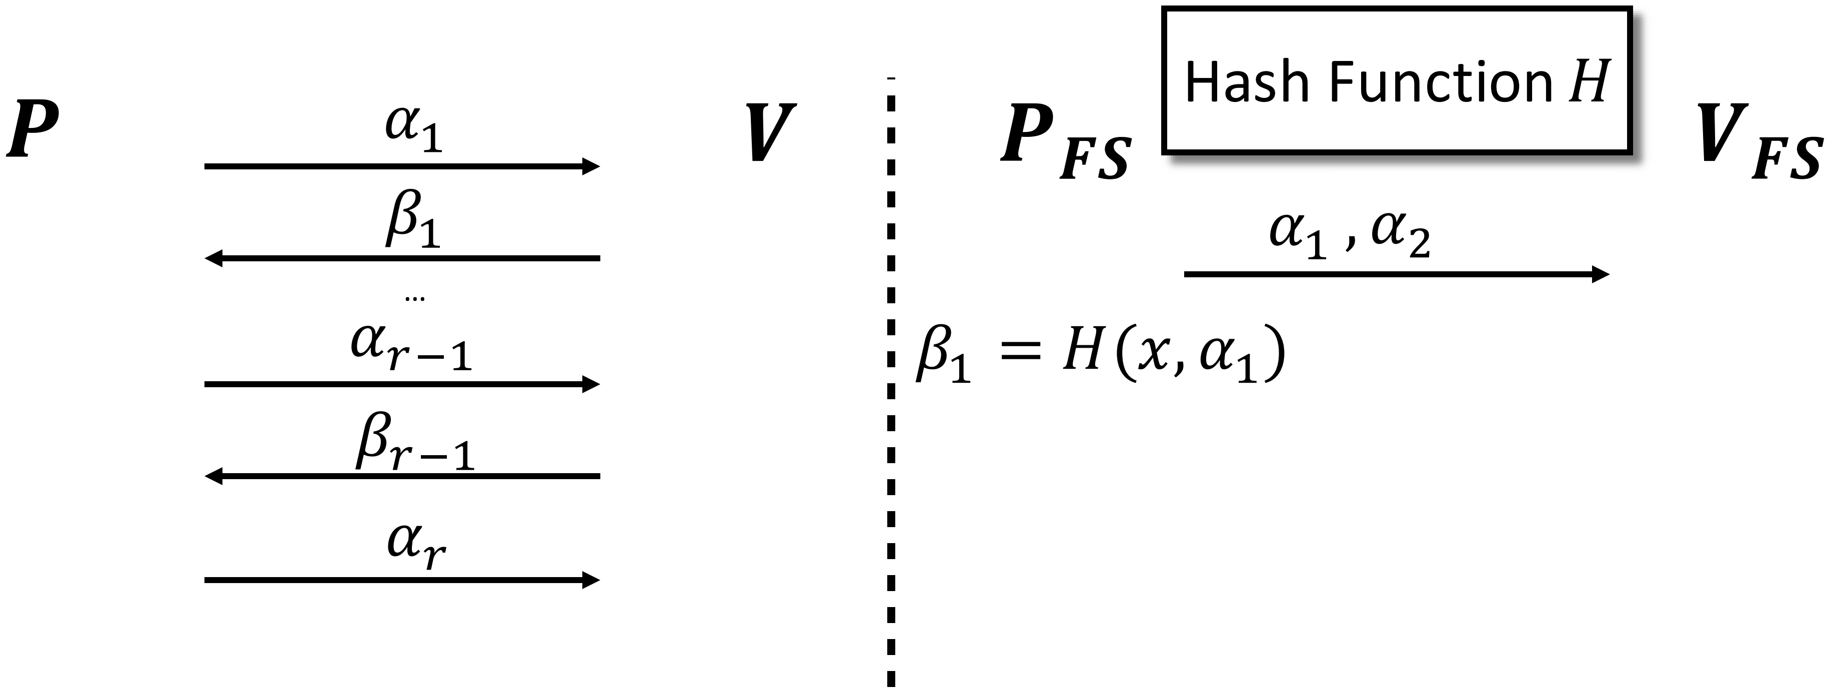

(Each β_i uniformly random)

The Fiat-Shamir Transform

Public-Coin Interactive Argument → generically → Non-Interactive Argument

$\boldsymbol{P}$ $\boldsymbol{V}$

α_1

β_1

…

α_{r-1}

β_{r-1}

α_r

(Each β_i uniformly random)

$\boldsymbol{P_{FS}}$ Hash Function H $\boldsymbol{V_{FS}}$

$\alpha_1, \ldots, \alpha_r$

$\beta_1 = H(x, \alpha_1)$

$\beta_2 = H(x, \alpha_1, \alpha_2)$

…

$\beta_i = H(x, \alpha_1, \ldots, \alpha_i)$

The Fiat-Shamir Transform

Extremely influential methodology.

Powerful: We know that interaction buys a lot. FS makes interaction free.

Practical: Very low overhead.

Expressive: Efficient Signature, CS proofs, (zk-)SNARGs, STARKs…

Fiat Shamir – Security?

Central question in cryptography:

Do there exist hash functions for which the Fiat-Shamir transform is secure?

Answer: we don't (quite) know ☹.

Still, would like to understand and so we'll analyze security assuming an *ideal* hash function.

The Random Oracle Model [BR93]

The random oracle model simply means that all parties are given blackbox access to a fully random function $R: \{0,1\}^{\lambda} \rightarrow \{0,1\}^{\lambda}$.

Security should hold whp over the choice of R.

Q: How should we view protocols secure in ROM?

A: TBD.

FS in the ROM

Public-Coin Interactive Argument → (generically) → Non-Interactive Argument

$\boldsymbol{P}$ — $\boldsymbol{V}$

α_1 →

β_1 ←

…

α_{r-1} →

β_{r-1} ←

α_r →

(Each β_i uniformly random)

$\boldsymbol{P_{FS}}$ — Random Oracle R — $\boldsymbol{V_{FS}}$

$\alpha_1, \dots, \alpha_r$ →

$\beta_1 = R(x, \alpha_1)$

$\beta_2 = R(x, \alpha_1, \alpha_2)$

…

$\beta_i = R(x, \alpha_1, \dots, \alpha_i)$

FS in the ROM

Thm [PS96,Folklore]: for every constant-round interactive argument Π with negl. soundness, whp over R, the protocol Π_R is secure.

Tightness

Claim: $\exists$multi-round protocol Π with negl. soundness error s.t. Π_{FS} is ***not*** sound (even in ROM).

Proof: Take any constant-round protocol with constant soundness and repeat sequentially.

Tightness

Public-Coin Interactive Argument

$\boldsymbol{P}$ $\quad$ $\boldsymbol{V}$

α_1 →

β_1 ←

…

α_{r-1} →

β_{r-1} ←

α_r →

Non-Interactive Argument

$\boldsymbol{P_{FS}}$ $\quad$ Hash Function H $\quad$ $\boldsymbol{V_{FS}}$

$\alpha_1, \ldots, \alpha_r$ →

$$\beta_1 = H(x, \alpha_1)$$

$$\beta_2 = H(x, \alpha_1, \alpha_2)$$

…

$$\beta_i = H(x, \alpha_1, \ldots, \alpha_i)$$

FS in the ROM

Thm [PS96,Folklore]: for every constant-round interactive argument Π with negl. soundness, whp over R, the protocol Π_R is secure.

(Actually extends to **some** multi-round protocols.)

We will see the proof in detail, but for simplicity focus on 3-message protocol.

FS in ROM

Public-Coin Interactive Protocol

Non-Interactive Argument

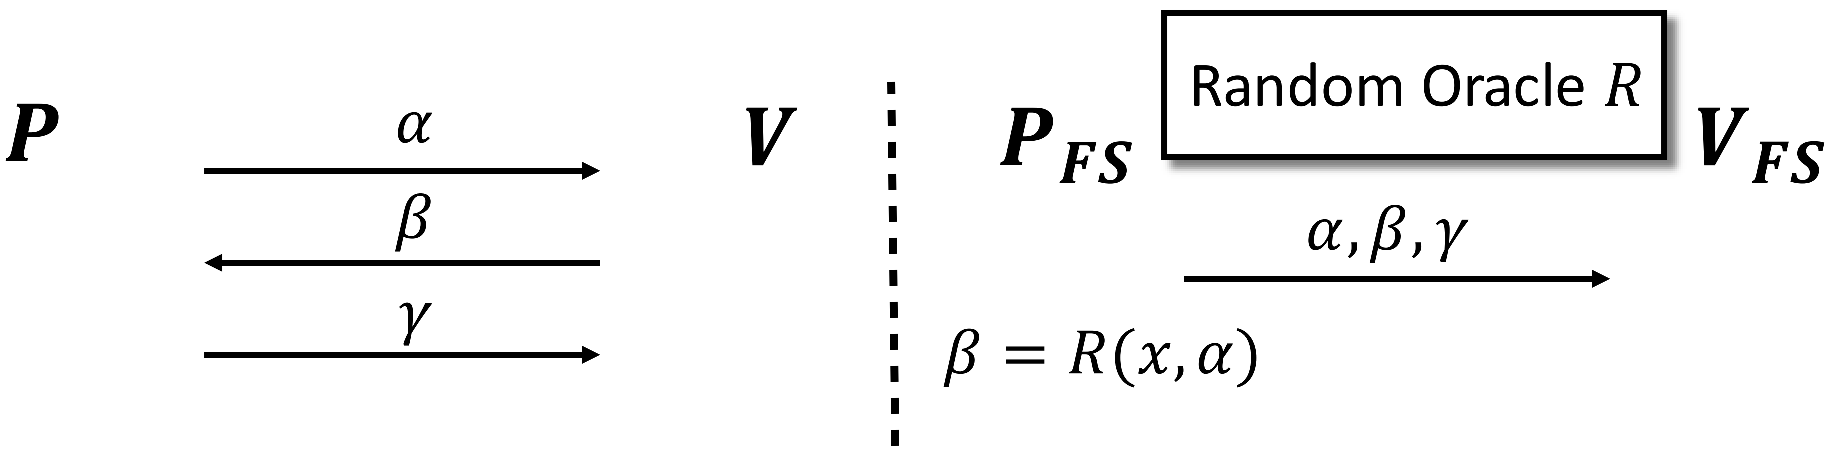

FS in ROM

Need to show:

- Completeness. 
- Soundness.
- Zero knowledge.

FS in ROM: Soundness

Suppose $\exists x \notin L$ and P_{FS}^{*} that runs in time T and makes V_{FS} accept x wp $\geq \epsilon$.

Will construct P^{*} s.t. V accepts x w.p. $\mathrm{poly}\left(\epsilon, \frac{1}{T}\right)$.

First, a Useful Fact

Fact: suppose (X, Y) are jointly distributed RVs s.t.

$$\Pr[A(X, Y)] \geq \epsilon.$$

Then, for at least $\epsilon/2$ fraction of $x's$ it holds that

$$(*) \Pr_{Y|x}[A(x, Y)] \geq \epsilon/2.$$

Proof: Markov's inequality.

First, a Useful Fact

Fact: suppose (X, Y) are jointly distributed RVs s.t.

$$\Pr[A(X, Y)] \geq \epsilon.$$

Then, for at least $\epsilon/2$ fraction of $x's$ it holds that

$$(*) \Pr_{Y|x}[A(x, Y)] \geq \epsilon/2.$$

Proof:

First, a Useful Fact

Fact: suppose (X, Y) are jointly distributed RVs s.t.

$$\Pr[A(X, Y)] \geq \epsilon.$$

Then, for at least $\epsilon/2$ fraction of $x's$ it holds that

$$(*) \Pr_{Y|x}[A(x, Y)] \geq \epsilon/2.$$

Proof: suppose not. Call x good if (*) holds

$$\Pr[A(X, Y)] =$$

First, a Useful Fact

Fact: suppose (X, Y) are jointly distributed RVs s.t.

$$\Pr[A(X,Y)] \geq \epsilon.$$

Then, for at least $\epsilon/2$ fraction of $x's$ it holds that

$$(*) \Pr_{Y|x}[A(x,Y)] \geq \epsilon/2.$$

Proof: suppose not. Call x good if (*) holds

$$\Pr[A(X,Y)] = \Pr[X \text{ good}] \cdot \Pr[A(X,Y)|X\ good] + \\ \Pr[X \text{ bad}] \cdot \Pr[A(X,Y)|X\ bad]$$

First, a Useful Fact

Fact: suppose (X, Y) are jointly distributed RVs s.t.

$$\Pr[A(X,Y)] \geq \epsilon.$$

Then, for at least $\epsilon/2$ fraction of $x's$ it holds that

$$(*) \Pr_{Y|x}[A(x,Y)] \geq \epsilon/2.$$

Proof: suppose not. Call x good if (*) holds

$$\Pr[A(X,Y)] = \Pr[X \text{ good}] \cdot \Pr[A(X,Y)|X\ good] + \Pr[X \text{ bad}] \cdot \Pr[A(X,Y)|X\ bad]$$

$$< \frac{\epsilon}{2} \cdot 1 + 1 \cdot \frac{\epsilon}{2}$$

First, a Useful Fact

Fact: suppose (X, Y) are jointly distributed RVs s.t.

$$\Pr[A(X,Y)] \geq \epsilon.$$

Then, for at least $\epsilon/2$ fraction of $x's$ it holds that

$$(*)\ \Pr_{Y|x}[A(x,Y)] \geq \epsilon/2.$$

Proof: suppose not. Call x good if (*) holds

$$\Pr[A(X,Y)] = \Pr[X \text{ good}] \cdot \Pr[A(X,Y)|X\ good] + \Pr[X \text{ bad}] \cdot \Pr[A(X,Y)|X\ bad]$$

$$< \frac{\epsilon}{2} \cdot 1 + 1 \cdot \frac{\epsilon}{2}$$

$$= \epsilon$$

FS in ROM: Soundness

Suppose $\exists x \notin L$ and P^*_{FS} that runs in time T and makes V_{FS} accept x wp $\geq \epsilon$.

Will construct P^* s.t. V accept x w.p. $\mathrm{poly}\left(\epsilon, \frac{1}{T}\right)$.

Soundness Analysis

Denote oracle queries by $Q_1, \ldots, Q_T$.

Wlog all Q_i's distinct and $\alpha \in \{Q_1, \ldots, Q_T\}$.

Claim: $\exists i^* \in [T]$ s.t. P^*_{FS} wins w.p. ϵ/T conditioned on $Q_{i^*} = \alpha$.

Proof: by contradiction.

“The Forking Lemma”

Key Lemma: for $\frac{\epsilon}{2T}$ fraction of $(q_1, \ldots, q_{i^*})$ it holds that P_{FS}^* wins w.p. $\frac{\epsilon}{2T}$ conditioned on $Q_{i^*} = \alpha$ and $Q_i = q_i$ for all $i \leq i^*$.

Proof: by useful fact.

Breaking Soundness of V

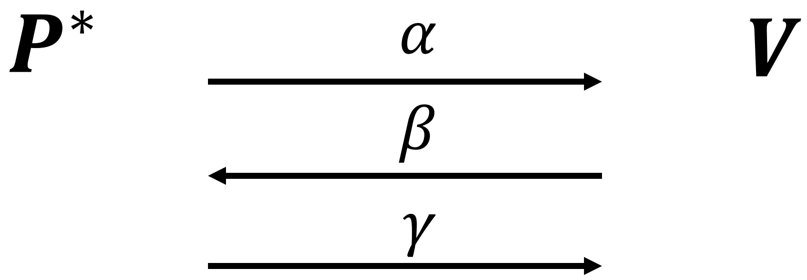

1. Start running P_{FS}^{*} up to it's i^{*}th query, using random answers.
2. Let $\alpha = Q_{i^{*}}$ be the i^{*}th query. Send α (and get β).
3. Continue running P_{FS}^{*} while answering $Q_{i^{*}}$ with β and other queries uniformly at random.
4. Eventually P_{FS}^{*} outputs $(\alpha', \beta', \gamma')$.
5. If $\alpha = \alpha'$ and $\beta = \beta'$ send $\gamma = \gamma'$.

Breaking Soundness of V: Analysis

Rely on forking lemma:

Forking Lemma: for $\frac{\epsilon}{2T}$ fraction of $(q_1, \ldots, q_{i^*})$ it holds that P_{FS}^* wins w.p. $\frac{\epsilon}{2T}$ conditioned on $Q_{i^*} = \alpha$ and $Q_i = q_i$ for all $i \leq i^*$.

Get that wp $\frac{\epsilon}{2T}$ over choice of $(Q_1, \ldots, Q_i)$ it holds that wp $\frac{\epsilon}{2T}$ over all remaining coin tosses that P_{FS}^* wins and $\alpha' = \alpha$.

$\Rightarrow$ our P^* wins wp $\left(\frac{\epsilon}{2T}\right)^2$, which is non-negligible.

FS in ROM: ZK

Have not defined ZK in the ROM and as there are multiple definitions (and issues).

Intuitively though, beyond seeing (α, β, γ) (which can be generated from x by (HV)-ZK), the verifier has obtained oracle access to a random function R such that $R(x, \alpha) = \beta$.

Could it have obtained such a function by itself?

Short answer: kind of…

Long answer: depends on the definition. ☺

FS in ROM

Conclusion: FS is sound in ROM (and ZK for some suitable definitions).

But we cannot use hash functions that take 2^{λ} bits to describe!

So, is the Fiat-Shamir transform secure?

Bad news [CHG98]: ∃ protocols secure in ROM but totally broken using *any* instantiation.

Fiat Shamir – Security?

Given negative result, how to interpret ROM proof of security?

Optimist's view:

- Counterexamples are contrived.
- ROM analysis $\Rightarrow$ strong indication FS is secure in real-life.
- ROM analysis = good heuristic. Can help both in terms of feasibility and efficiency.

Pessimist's view:

- Basing security on an assumption that we do not understand, and have a negative indication for, is undesirable if not flat out dangerous.

Instantiating Fiat Shamir with Explicit Hash function

A Basic Question

Can we instantiate the heuristic securely using an explicit hash family?

<u>**Def:**</u> a hash family H is FS-compatible for a Π if $FS_H(\Pi)$ is "secure".

$\boldsymbol{P_{FS}}$ $\qquad\qquad$ $\boldsymbol{V_{FS}}$

$\xleftarrow{\quad h \quad}$ $\quad h \in H$

$\beta = h(x, \alpha)$ $\quad \xrightarrow{\quad \alpha, \beta, \gamma \quad}$

FS using Explicit Family

Need to consider soundness & zero-knowledge.

Start with zero-knowledge.

Def: H is programmable if can sample random $h \in H$ conditioned on $h(x, \alpha) = \beta$.

ZK for FS

Claim: if H is programmable and Π is HVZK $\Rightarrow \Pi_{FS}(h)$ is ZK.

Proof: construct simulator.

1. Sample (α, β, γ).
2. Sample H conditioned on $H(x, \alpha) = \beta$.
3. Output $\big(H, (\alpha, \beta, \gamma)\big)$.

Exercise: show dist. identical.

Soundness for FS

Thm [B01,GK03]: ∃ protocols which are not FS-compatible for any H.

Hope? Those counterexamples are arguments! Maybe sound if we start with a proof?

[BDGJKLW13]: no blackbox reduction to a falsifiable assumption, **even for proofs**.

Fiat Shamir for Proofs?

- Stay tuned for afternoon talk.

- Closely related to the question of parallel composition of ZK [DNRS03].

Thanks!