Witness Semantic Security

Paul Lou[†], Nathan Manohar[‡], Amit Sahai[†]

[†]UCLA, Los Angeles, CA

[‡]IBM T.J. Watson Research Center, Yorktown Heights, NY

Eurocrypt 2024

Two-round Publicly-verifiable Setting

(Babai '85, Goldwasser, Sipser ' 86, Fortnow '87, Aiello, Hastad '87, Goldreich, Oren '94)

$$x \in L \in \text{NP}$$

"I swear on Merlin's beard that x is in L."

"Convince me! I want mathematical proof, not witchcraft."

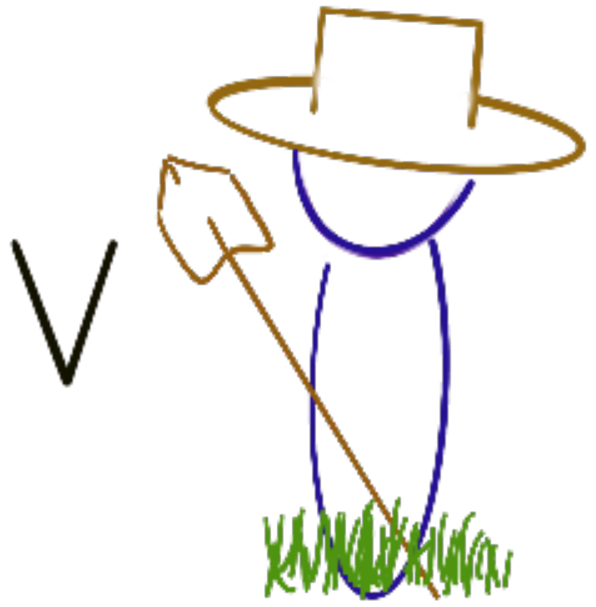

Public verifiability: Anyone (who trusts the Verifier) can use the first round message to verify the second round message!

- Implied by public-coin (i.e. Arthur-Merlin [AM] protocols).
- Typically allows the first message to be *reused for multiple proofs!*

What kind of security can we guarantee?

General Cryptographic Proof Systems for **NP**

(Goldwasser, Micali, Rackoff '85, Goldreich, Micali, Widgerson, '86)

$x \in L \in \mathsf{NP}$

"I swear on Merlin's beard that x is in L."

"Convince me! I want mathematical proof, not witchcraft."

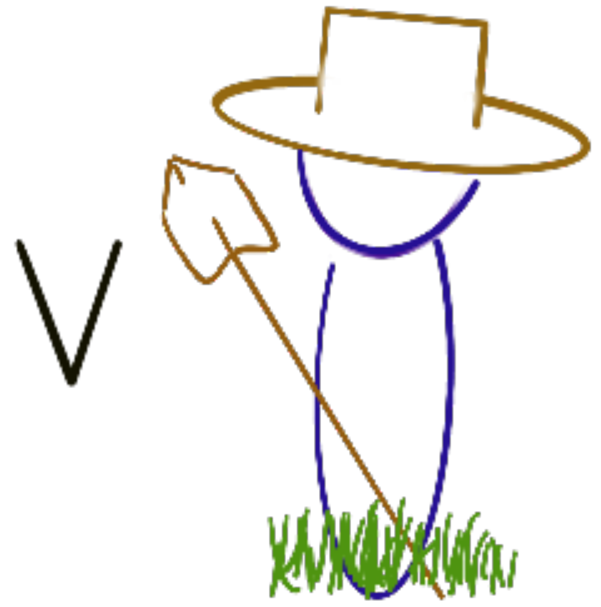

Security (Zero-knowledge): Convinced but doesn't know more than the validity of the statement.

Goldreich, Oren '94, Barak, Lindell, Vadhan '04: At least three rounds of messaging is necessary for ZK.

Two-round Publicly-verifiable Setting

(Babai '85, Goldwasser, Sipser ' 86, Fortnow '87, Aiello, Hastad '87, Goldreich, Oren '94)

What kind of security can we guarantee?

- ☑ Witness indistinguishability (WI) (Feige, Shamir 1990; Dwork, Naor 2000; Groth, Ostrovsky, Sahai 2006)
- ☑ Witness hiding (WH) (Feige, Shamir 1990; Pass 2003; Bitansky, Khurana, Paneth 2019; Kuykendall, Zhandry 2020)
- ☑ Super-polynomial simulation (SPS) (Pass 2003)

Two-round Publicly-verifiable Setting

(Babai '85, Goldwasser, Sipser ' 86, Fortnow '87, Aiello, Hastad '87, Goldreich, Oren '94)

What kind of security can we guarantee?

- ☑ Witness indistinguishability (WI) (Feige, Shamir 1990; Dwork, Naor 2000; Groth, Ostrovsky, Sahai 2006)
- ☑ Witness hiding (WH) (Feige, Shamir 1990; Pass 2003; Bitansky, Khurana, Paneth 2019; Kuykendall, Zhandry 2020)
- ☑ Super-polynomial simulation (SPS) (Pass 2003)

What is the qualitative security guarantee?

Consider an encrypted signed document with three sensitive fields of information, e.g. social security number or month-by-month financial transactions.

Two-round Publicly-verifiable Setting

(Babai '85, Goldwasser, Sipser ' 86, Fortnow '87, Aiello, Hastad '87, Goldreich, Oren '94)

~~What kind of security can we guarantee?~~

- ☑ Witness indistinguishability (WI) (Feige, Shamir 1990; Dwork, Naor 2000; Groth, Ostrovsky, Sahai 2006)
- ☑ Witness hiding (WH) (Feige, Shamir 1990; Pass 2003; Bitansky, Khurana, Paneth 2019; Kuykendall, Zhandry 2020)
- ☑ Super-polynomial simulation (SPS) (Pass 2003)

What is the qualitative security guarantee?

Consider an encrypted signed document with three sensitive fields of information, e.g. social security number or month-by-month financial transactions.

- WI: meaningless if the encryption scheme has perfect correctness, i.e. unique witness :(
- WH: doesn't prevent partial information loss :(
- SPS: leaks information computable in super-polynomial time, not easy to interpret :(

Two-round Publicly-verifiable Setting

(Babai '85, Goldwasser, Sipser ' 86, Fortnow '87, Aiello, Hastad '87, Goldreich, Oren '94)

What kind of security can we guarantee?

- Witness indistinguishability (WI) (Feige, Shamir 1990; Dwork, Naor 2000; Groth, Ostrovsky, Sahai 2006)
- Witness hiding (WH) (Feige, Shamir 1990; Pass 2003; Bitansky, Khurana, Paneth 2019; Kuykendall, Zhandry 2020)
- Super-polynomial simulation (SPS) (Pass 2003)

Can we have stronger qualitative guarantees?

Two-round Publicly-verifiable Setting

(Babai '85, Goldwasser, Sipser ' 86, Fortnow '87, Aiello, Hastad '87, Goldreich, Oren '94)

What kind of security can we guarantee?

- ☑ Witness indistinguishability (WI) (Feige, Shamir 1990; Dwork, Naor 2000; Groth, Ostrovsky, Sahai 2006)
- ☑ Witness hiding (WH) (Feige, Shamir 1990; Pass 2003; Bitansky, Khurana, Paneth 2019; Kuykendall, Zhandry 2020)
- ☑ Super-polynomial simulation (SPS) (Pass 2003)

Can we have stronger qualitative guarantees?

Goldreich, Oren '94, Barak, Lindell, Vadhan '04: At least three rounds of messaging is necessary for ZK.

Two-round Publicly-verifiable Setting

(Babai '85, Goldwasser, Sipser ' 86, Fortnow '87, Aiello, Hastad '87, Goldreich, Oren '94)

~~**What kind of security can we guarantee?**~~

- ☑ Witness indistinguishability (WI) (Feige, Shamir 1990; Dwork, Naor 2000; Groth, Ostrovsky, Sahai 2006)
- ☑ Witness hiding (WH) (Feige, Shamir 1990; Pass 2003; Bitansky, Khurana, Paneth 2019; Kuykendall, Zhandry 2020)
- ☑ Super-polynomial simulation (SPS) (Pass 2003)

Can we have stronger qualitative guarantees?

Goldreich, Oren '94 (as noted by Bitansky, Khurana, Paneth '19):
Even *weak zero-knowledge* (Dwork, Naor, Reingold, Stockmeyer '03)) is *impossible* in the two-round publicly-verifiable setting!

Two-round Publicly-verifiable Setting

(Babai '85, Goldwasser, Sipser ' 86, Fortnow '87, Aiello, Hastad '87, Goldreich, Oren '94)

~~What kind of security can we guarantee?~~

- ☑ Witness indistinguishability (WI) (Feige, Shamir 1990; Dwork, Naor 2000; Groth, Ostrovsky, Sahai 2006)
- ☑ Witness hiding (WH) (Feige, Shamir 1990; Pass 2003; Bitansky, Khurana, Paneth 2019; Kuykendall, Zhandry 2020)
- ☑ Super-polynomial simulation (SPS) (Pass 2003)

Can we have stronger qualitative guarantees?

There is a large gap in qualitative guarantees between the above and weak zero-knowledge.

Goldreich, Oren '94 (as noted by Bitansky, Khurana, Paneth '19):

Even *weak zero-knowledge* (Dwork, Naor, Reingold, Stockmeyer '03)) is *impossible* in the two-round publicly-verifiable setting!

Two-round Publicly-verifiable Setting

(Babai '85, Goldwasser, Sipser ' 86, Fortnow '87, Aiello, Hastad '87, Goldreich, Oren '94)

~~What kind of security can we guarantee?~~

- ☑ Witness indistinguishability (WI) (Feige, Shamir 1990; Dwork, Naor 2000; Groth, Ostrovsky, Sahai 2006)
- ☑ Witness hiding (WH) (Feige, Shamir 1990; Pass 2003; Bitansky, Khurana, Paneth 2019; Kuykendall, Zhandry 2020)
- ☑ Super-polynomial simulation (SPS) (Pass 2003)

Can we have stronger qualitative guarantees?

Yes! Addressing this gap...

In this work:

- We introduce the notion of ***Witness Semantic Security (WSS)***.
- We construct a two-round publicly-verifiable cryptographic argument satisfying WSS from the subexponential hardness of LWE.

Goldreich, Oren '94 (as noted by Bitansky, Khurana, Paneth '19):

Even *weak zero-knowledge* (Dwork, Naor, Reingold, Stockmeyer '03)) is *impossible* in the two-round publicly-verifiable setting!

Intuition: Witness Semantic Security (WSS)

* **Encryption semantic security** (Goldwasser, Micali '82): Information about the message that can be computed given the ciphertext can also be computed without the ciphertext.

* **Witness semantic security**: Information about the witness that can be computed given the proof can also be computed with only the statement.

Intuition: Witness Semantic Security (WSS)

* **Encryption semantic security** (Goldwasser, Micali '82): Information about the message that can be computed given the ciphertext can also be computed without the ciphertext.

* **Witness semantic security**: Information about the witness that can be computed given the proof can also be computed with only the statement.

A witness semantic secure proof hides all non-trivial partial information about the witness.

This Work: Witness Semantic Security (WSS)

Definition (basic variant): A two-round interactive argument system (P, V) for an NP language L is WSS if for all polynomially-bounded probability ensembles D over

$$\{(x, w, \mathsf{aux}, f, y) \mid y = f(w), (x, w) \in R_L, f \text{ deterministic}\}$$

for all polynomial sized A_1, A_2 there exists a polynomial sized B and a negligible function $\mu(\cdot)$ such that

$$\Pr\left[A_2(1^\lambda, x, f, \langle P(x, w), A_1(1^\lambda)\rangle, \mathsf{aux}) = y\right] \leq \Pr\left[B(1^\lambda, x, f, \mathsf{aux}) = y\right] + \mu(\lambda).$$

Definition is in the *delayed-input model* in the two-round setting, when the first round (honest & malicious) Verifier message is independent of the statement.

This Work: Witness Semantic Security (WSS)

Definition (basic variant): A two-round interactive argument system (P, V) for an NP language L is WSS if for all polynomially-bounded probability ensembles D over

$$\{(x, w, \mathsf{aux}, f, y) \mid y = f(w), (x, w) \in R_L, f \text{ deterministic}\}$$

for all polynomial sized A_1, A_2 there exists a polynomial sized B and a negligible function $\mu(\cdot)$ such that

$$\Pr\left[A_2(1^\lambda, x, f, \langle P(x, w), A_1(1^\lambda)\rangle, \mathsf{aux}) = y\right] \leq \Pr\left[B(1^\lambda, x, f, \mathsf{aux}) = y\right] + \mu(\lambda).$$

WSS morally looks like zero-knowledge!

This Work: Witness Semantic Security (WSS)

Definition (basic variant): A two-round interactive argument system (P, V) for an NP language L is WSS if for all polynomially-bounded probability ensembles D over

$$\{(x, w, \mathsf{aux}, f, y) \mid y = f(w), (x, w) \in R_L, f \text{ deterministic}\}$$

for all polynomial sized A_1, A_2 there exists a polynomial sized B and a negligible function $\mu(\cdot)$ such that

$$\Pr\left[A_2(1^\lambda, x, f, \langle P(x, w), A_1(1^\lambda)\rangle, \mathsf{aux}) = y\right] \leq \Pr\left[B(1^\lambda, x, f, \mathsf{aux}) = y\right] + \mu(\lambda).$$

WSS morally looks like zero-knowledge!

So why does this definition not imply distributional ZK?

This Work: Witness Semantic Security (WSS)

Definition (basic variant): A two-round interactive argument system (P, V) for an NP language L is WSS if for all polynomially-bounded probability ensembles D over

$$\{(x, w, \mathsf{aux}, f, y) \mid y = f(w), (x, w) \in R_L, f \text{ deterministic}\}$$

for all polynomial sized A_1, A_2 there exists a polynomial sized B and a negligible function $\mu(\cdot)$ such that

$$\Pr\left[A_2(1^\lambda, x, f, \langle P(x, w), A_1(1^\lambda)\rangle, \mathsf{aux}) = y\right] \leq \Pr\left[B(1^\lambda, x, f, \mathsf{aux}) = y\right] + \mu(\lambda).$$

WSS morally looks like zero-knowledge!

So why does this definition not imply distributional ZK?

First observe that this definition only considers a specific witness w.

Verifiable Witness Semantic Secure (VWSS)

Definition [VWSS]: A two-round interactive argument system (P, V) for an NP language L is VWSS if for all polynomially-bounded probability ensembles D over

$$\{(x, w, \mathsf{aux}, f) \mid (x, w) \in R_L, f \text{ deterministic and verifiable input/output}\}$$

where aux contains $V_f(\cdot, \cdot)$ for all polynomial sized A_1, A_2 there exists a polynomial sized B and a negligible function $\mu(\cdot)$ such that

$$\Pr\left[A_2(1^\lambda, x, f, \langle P(x, w), A_1(1^\lambda)\rangle, \mathsf{aux}) = y : \exists \tilde{w}, y = f(\tilde{w}) \wedge (x, \tilde{w}) \in R_L\right]$$

$$\leq \Pr\left[B(1^\lambda, x, f, \mathsf{aux}) = y : \exists \tilde{w}, y = f(\tilde{w}) \wedge (x, \tilde{w}) \in R_L\right] + \mu(\lambda).$$

$$V_f(x, y) = 1 \iff \exists \tilde{w}, ((x, \tilde{w}) \in R_L) \wedge (f(\tilde{w}) = y)$$

Verifiable Witness Semantic Secure (VWSS)

Definition [VWSS]: A two-round interactive argument system (P, V) for an NP language L is VWSS if for all polynomially-bounded probability ensembles D over

$$\{(x, w, \mathsf{aux}, f) \mid (x, w) \in R_L, f \text{ deterministic and verifiable input/output}\}$$

where aux contains $V_f(\cdot,\cdot)$ for all polynomial sized A_1, A_2 there exists a polynomial sized B and a negligible function $\mu(\cdot)$ such that

$$\Pr\left[A_2(1^\lambda, x, f, \langle P(x, w), A_1(1^\lambda)\rangle, \mathsf{aux}) = y : \exists \tilde{w}, y = f(\tilde{w}) \wedge (x, \tilde{w}) \in R_L\right]$$

$$\leq \Pr\left[B(1^\lambda, x, f, \mathsf{aux}) = y : \exists \tilde{w}, y = f(\tilde{w}) \wedge (x, \tilde{w}) \in R_L\right] + \mu(\lambda).$$

VWSS also morally looks like zero-knowledge! So what's different?

Verifiable Witness Semantic Secure (VWSS)

Definition [VWSS]: A two-round interactive argument system (P, V) for an NP language L is VWSS if for all polynomially-bounded probability ensembles D over

$$\{(x, w, \mathsf{aux}, f) \mid (x, w) \in R_L, f \text{ deterministic and verifiable input/output}\}$$

where **aux** contains $V_f(\cdot, \cdot)$ for all polynomial sized A_1, A_2 there exists a polynomial sized B and a negligible function $\mu(\cdot)$ such that

$$\Pr\left[A_2(1^\lambda, x, f, \langle P(x, w), A_1(1^\lambda)\rangle, \mathsf{aux}) = y : \exists \tilde{w}, y = f(\tilde{w}) \wedge (x, \tilde{w}) \in R_L\right]$$

$$\leq \Pr\left[B(1^\lambda, x, f, \mathsf{aux}) = y : \exists \tilde{w}, y = f(\tilde{w}) \wedge (x, \tilde{w}) \in R_L\right] + \mu(\lambda).$$

VWSS also morally looks like zero-knowledge! So what's different?

Observation: Existing simulation-based definitions of ZK ensures the hiding of *all* non-trivial information of the transcript.

This prevents the Prover from revealing something non-trivial (possibly inefficiently computable) about the Verifier's first message that the Verifier itself does not know!!

WSS and VWSS **allows** this behavior (remember this, we'll revisit this)!

Witness Semantic Security (WSS)

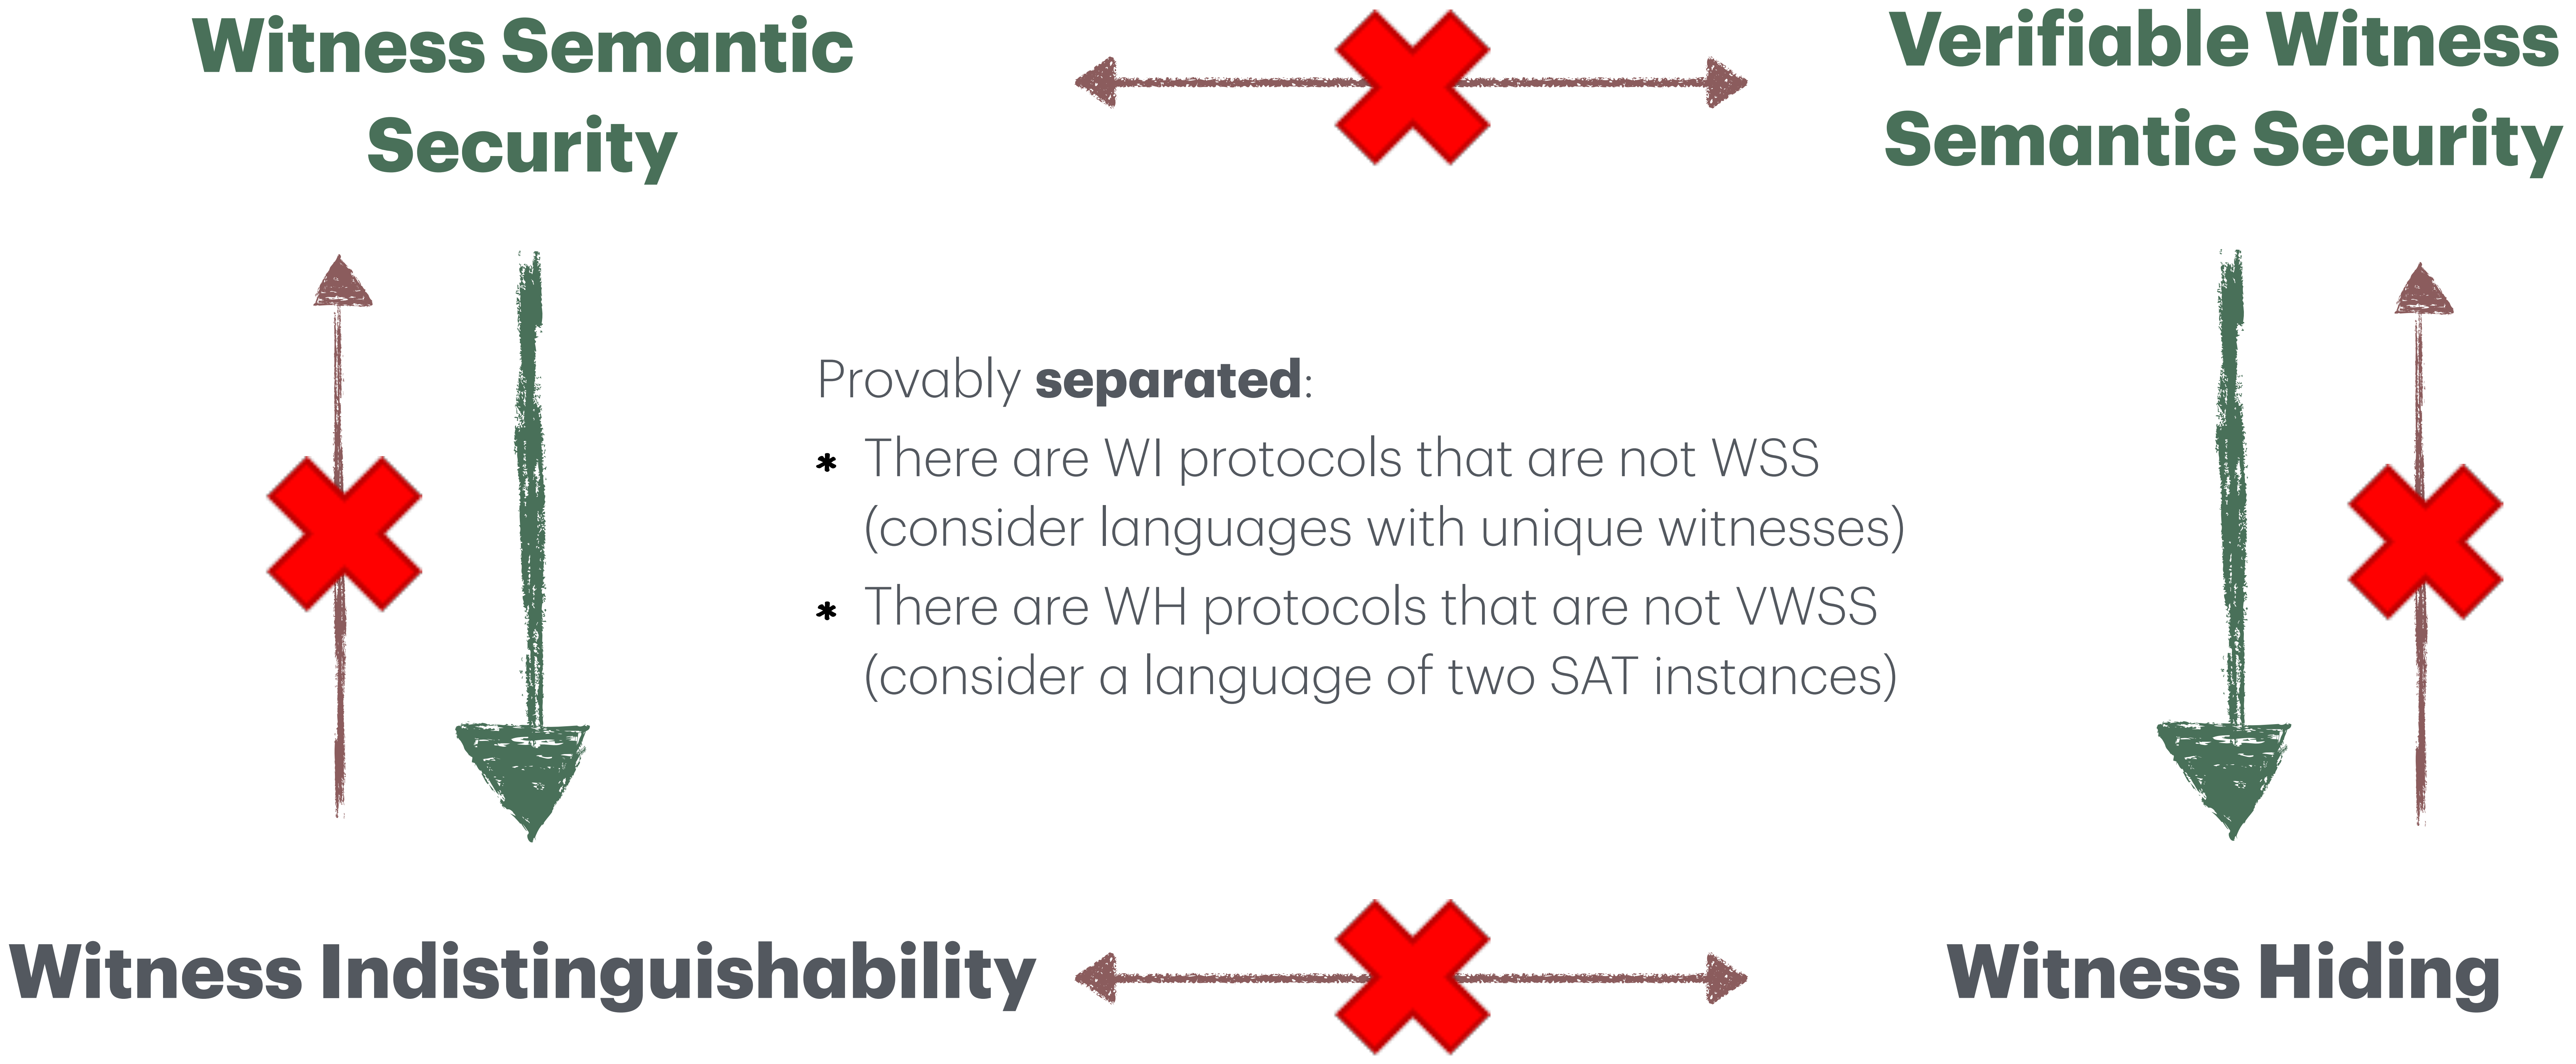

Provably **separated**:

* There are WI protocols that are not WSS (consider languages with unique witnesses)
* There are WH protocols that are not VWSS (consider a language of two SAT instances)

Witness Semantic Security (WSS)

We'll soon show a security notion that implies both!

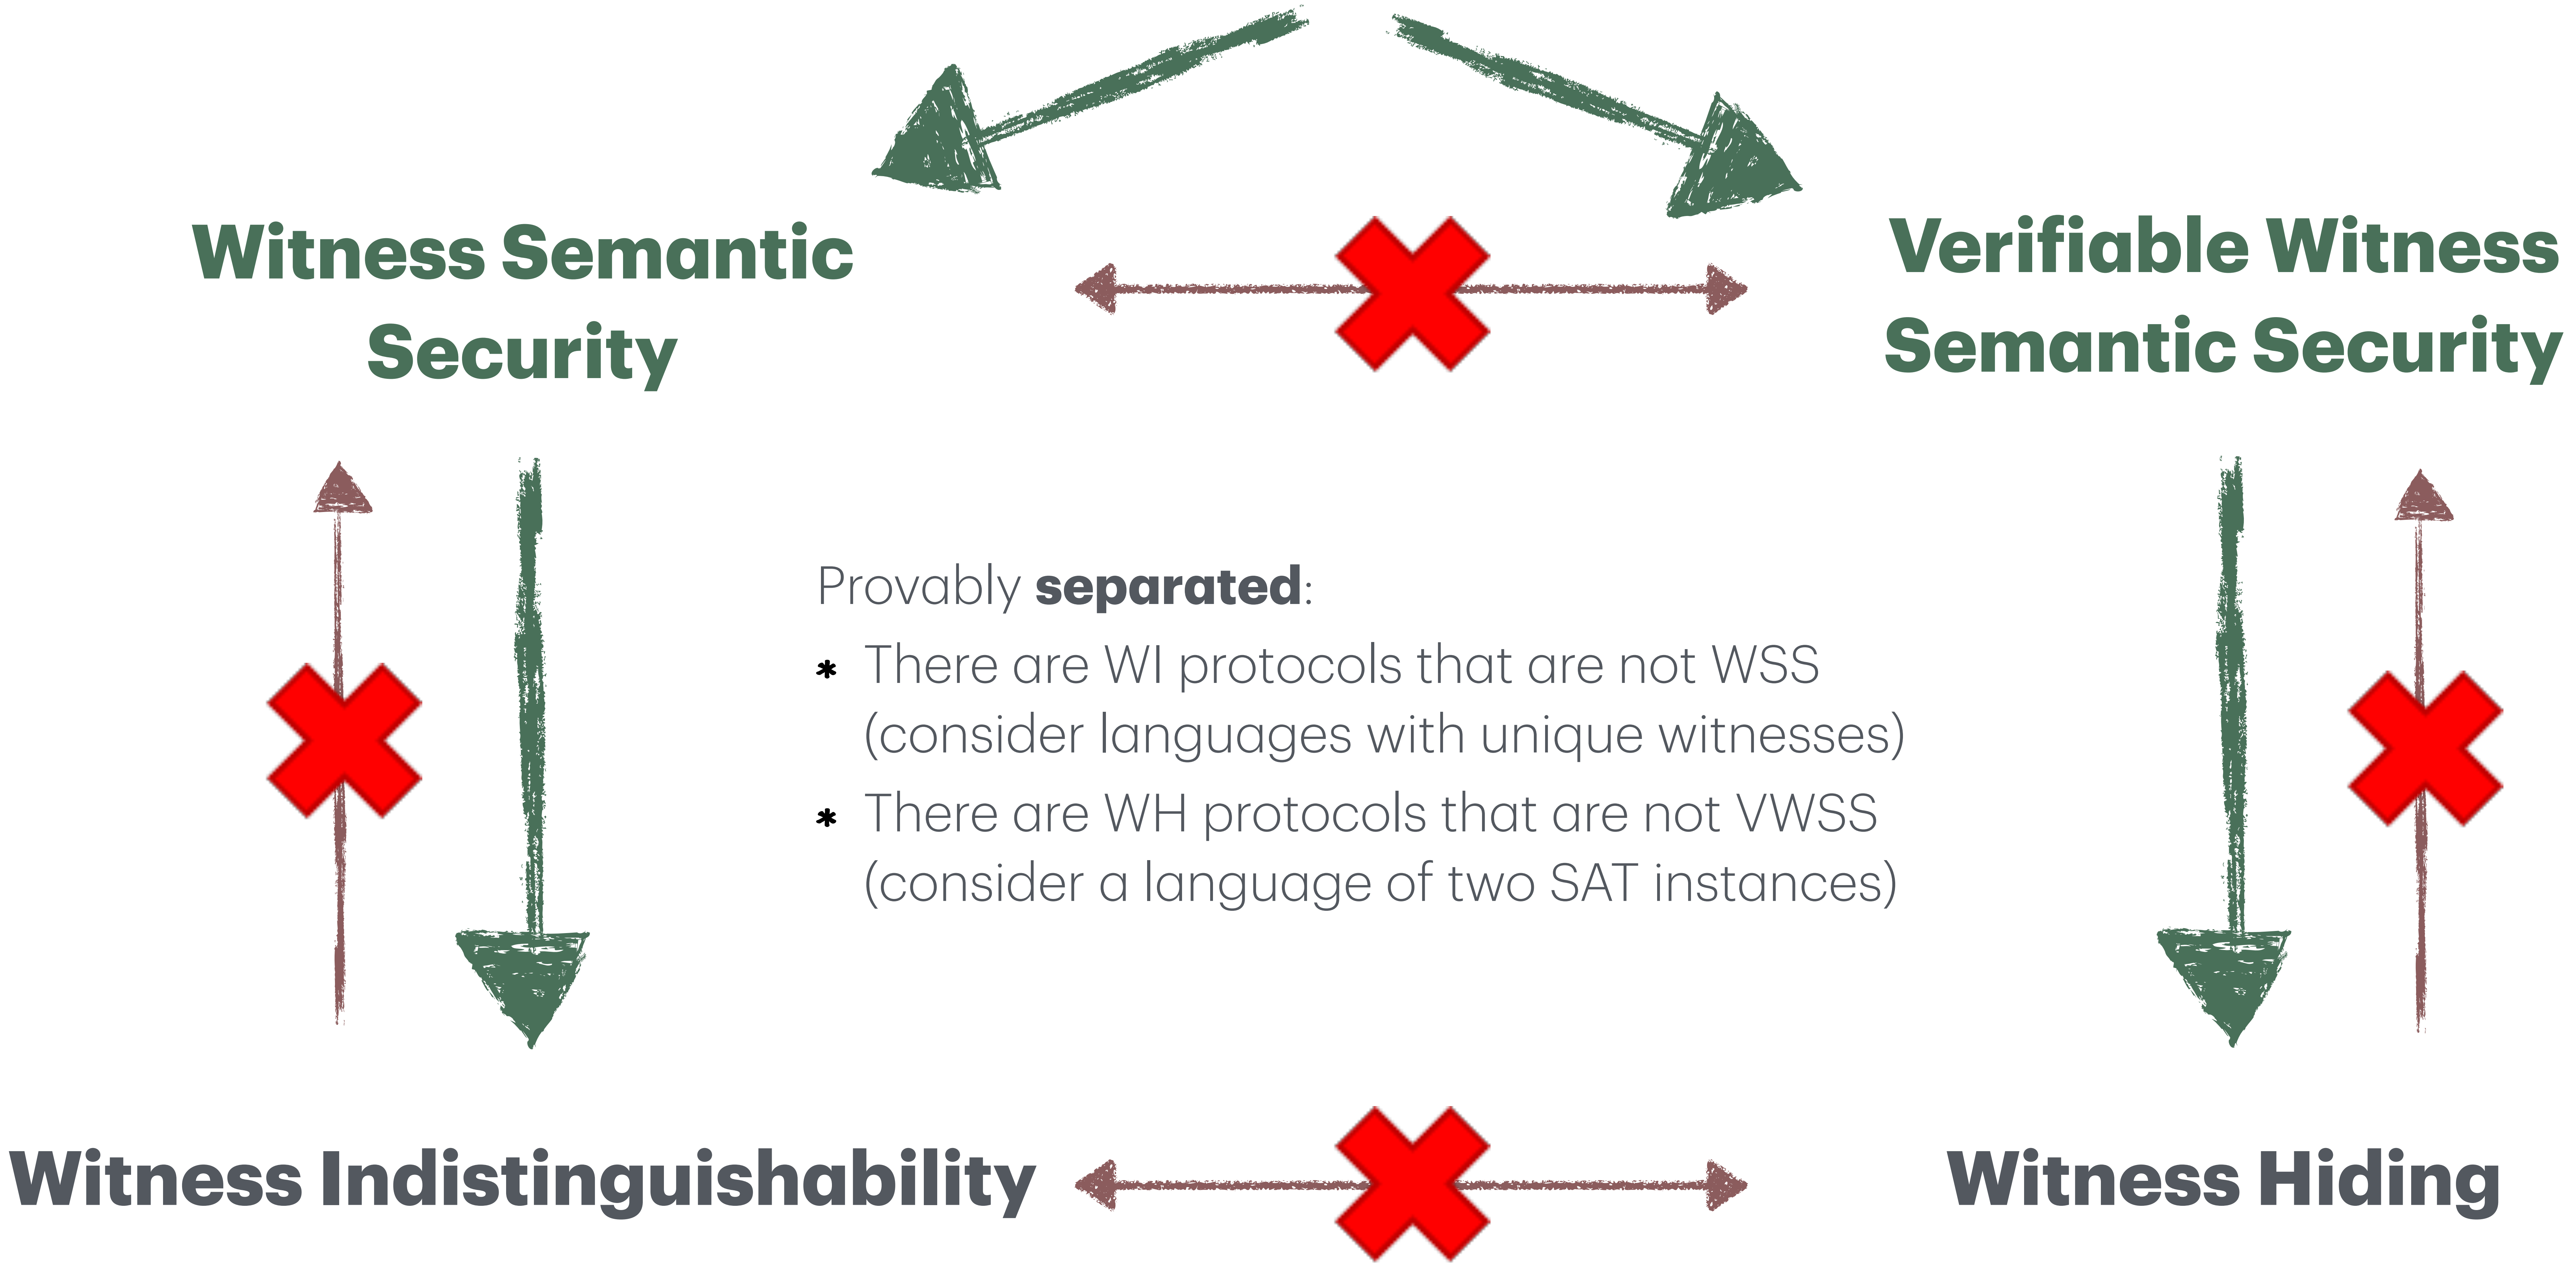

Provably **separated**:

* There are WI protocols that are not WSS (consider languages with unique witnesses)
* There are WH protocols that are not VWSS (consider a language of two SAT instances)

Another Viewpoint on Two-round Protocols: CRS-model Non-interactive Proof Systems

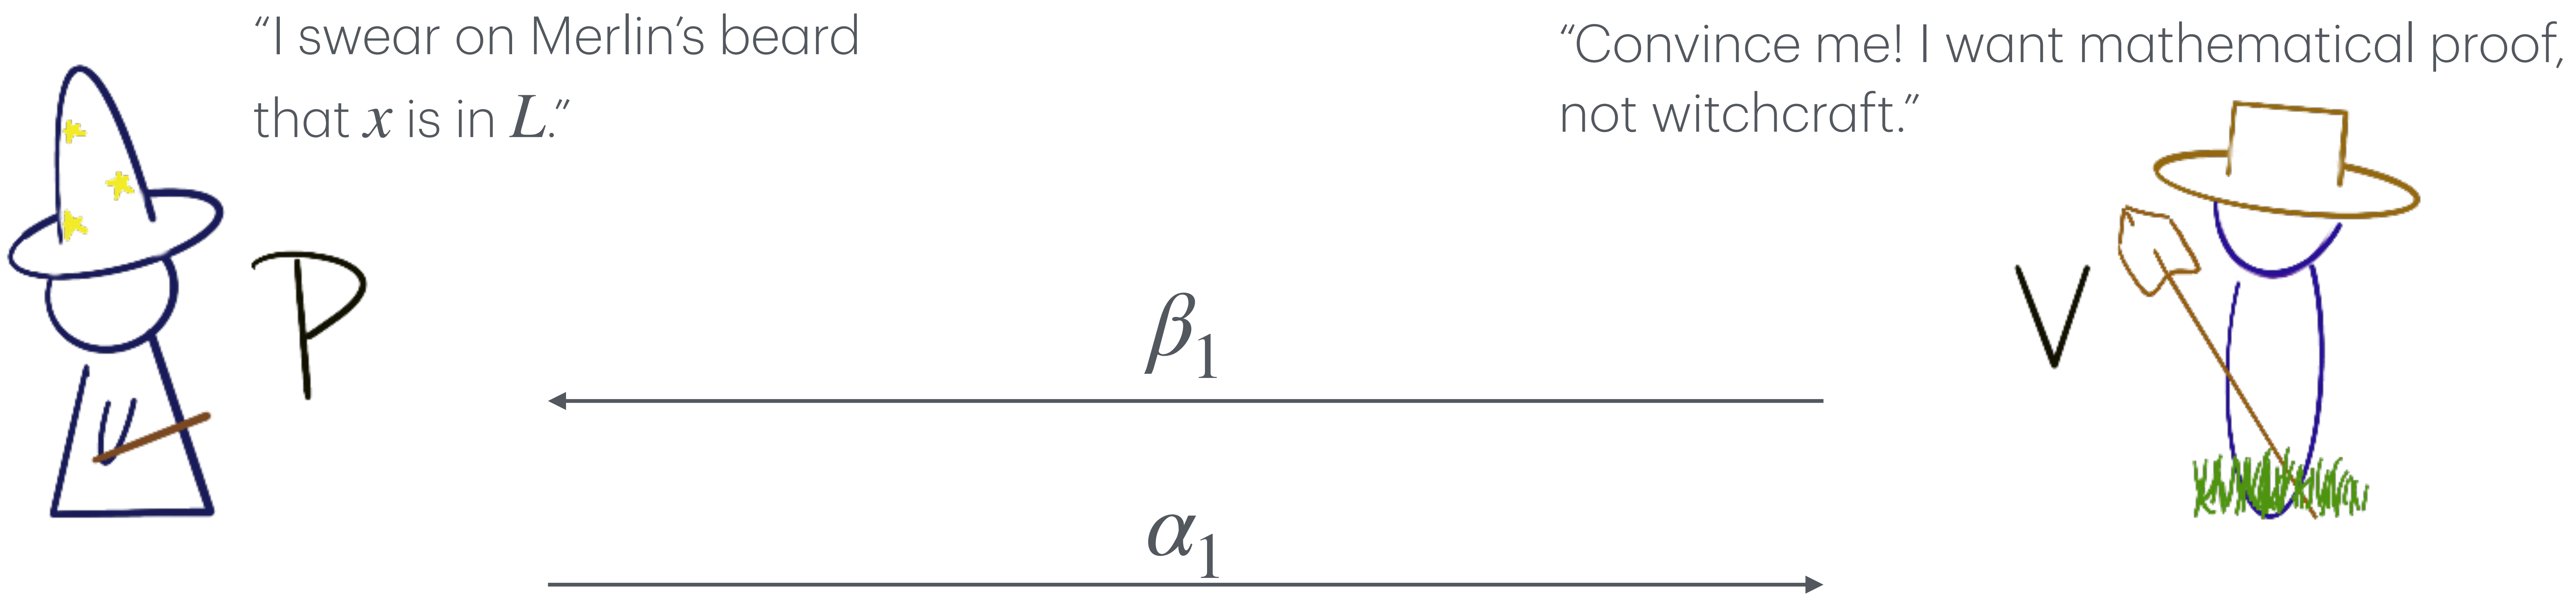

Another Viewpoint on Two-round Protocols: CRS-model Non-interactive Proof Systems

$\mathrm{CRS} \leftarrow \beta_1$

"I swear on Merlin's beard that x is in L."

"Convince me! I want mathematical proof, not witchcraft."

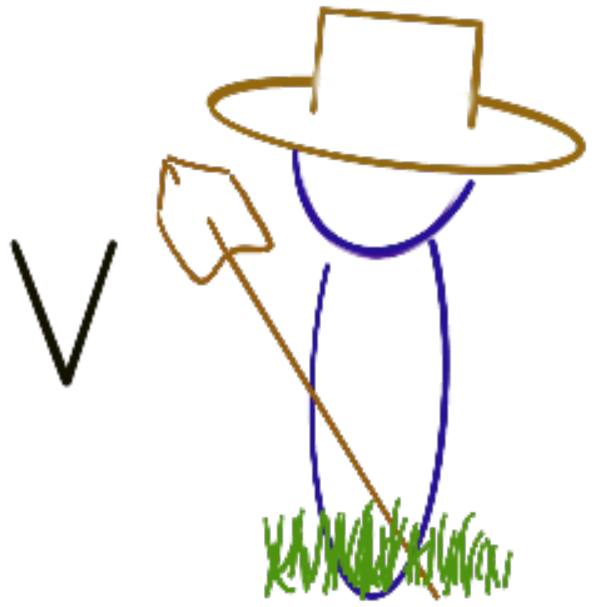

$$\xrightarrow{\quad \alpha_1 \quad}$$

Another Viewpoint on Two-round Protocols: CRS-model Non-interactive Proof Systems

$$\text{CRS} \leftarrow \beta_1$$

"I swear on Merlin's beard that x is in L."

"Convince me! I want mathematical proof, not witchcraft."

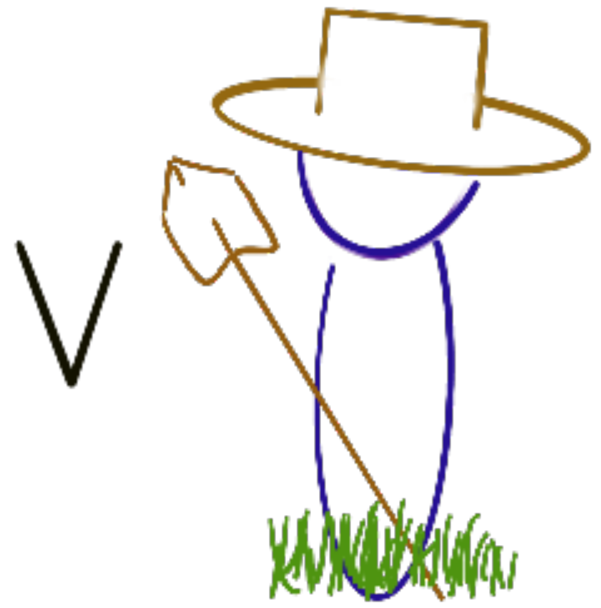

$$\xrightarrow{\alpha_1}$$

A key difference b/w standard 2-round and NIZK is that the CRS is statement independent.

Instead, this corresponds to the *delayed-input model* in the two-round setting, when the first round (honest & malicious) Verifier message is independent of the statement.

Natural Application of Two-round Protocols:
Malicious CRS Non-interactive Proof Systems

$\text{CRS} \leftarrow \beta_1$

"I swear on Merlin's beard that x is in L."

"Convince me! I want mathematical proof, not witchcraft."

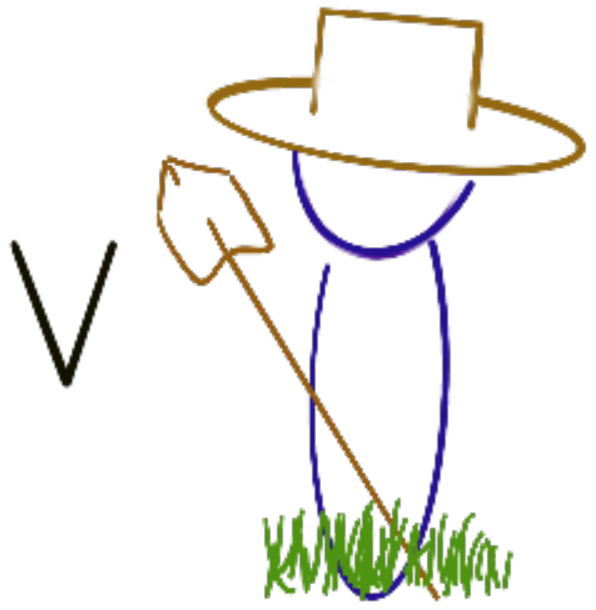

$\xrightarrow{\alpha_1}$

Even if the CRS is maliciously generated, the ZK* property of the two-round protocol preserves ZK* against a malicious V (no guarantees on soundness).

Natural Application of Two-round Protocols: Malicious CRS Non-interactive Proof Systems

$\text{CRS} \leftarrow \beta_1$

"I swear on Merlin's beard that x is in L."

"Convince me! I want mathematical proof, not witchcraft."

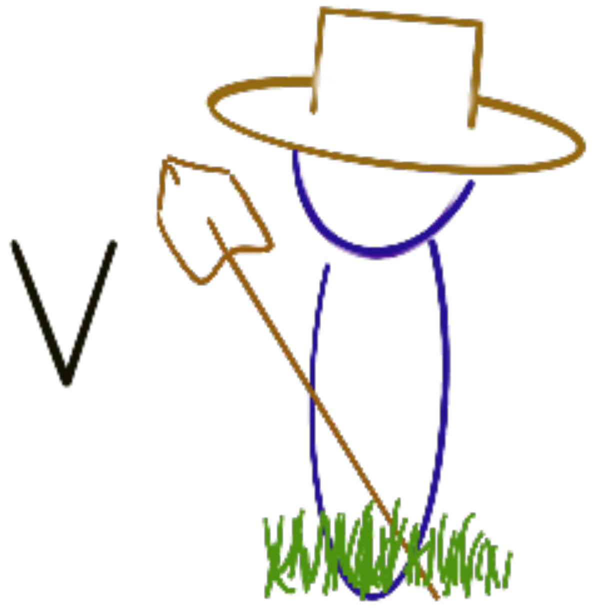

$\xrightarrow{\alpha_1}$

Even if the CRS is maliciously generated, the ZK* property of the two-round protocol preserves ZK* against a malicious V (no guarantees on soundness).

Bellare, Fuchsbauer, Scafuro '16: If soundness holds in the malicious CRS setting, then zero-knowledge cannot hold even in the *honest* CRS setting.

This Work: New Notion of Simulation (NUZK)

Definition (Standard Non-interactive Zero-Knowledge): There exists a PPT algorithm (S_1, S_2) such that for all PPT adversaries $\mathscr{A}$, the following is indistinguishable to the real world:

1. $\mathsf{CRS}, \tau \leftarrow S_1(1^\lambda)$.
2. $(x, w) \leftarrow \mathscr{A}(1^\lambda, \mathsf{CRS}), (x, w) \in R_L$.
3. $\pi \leftarrow S_2(x, \tau)$.

Definition (Non-Uniform Zero-Knowledge [NUZK] with Auxiliary Input): The simulator now depends non-uniformly on the CRS. For all CRS, there exists a circuit S_{CRS}, such that for all (x, w, Aux),

$$(x, \mathsf{CRS}, \mathsf{Prove}(\mathsf{CRS}, x, w), \mathsf{Aux}) \approx_c (x, \mathsf{CRS}, S_{\mathsf{CRS}}(x, \mathsf{Aux}), \mathsf{Aux})$$

This Work: New Notion of Simulation (NUZK)

Definition (Non-Uniform Zero-Knowledge [NUZK] with Auxiliary Input): The simulator now depends non-uniformly on the CRS. For all CRS, there exists a circuit S_{CRS}, such that for all (x, w, Aux),

$$(x, \mathsf{CRS}, \mathsf{Prove}(\mathsf{CRS}, x, w), \mathsf{Aux}) \approx_c (x, \mathsf{CRS}, S_{\mathsf{CRS}}(x, \mathsf{Aux}), \mathsf{Aux})$$

This Work: New Notion of Simulation (NUZK)

Definition (Non-Uniform Zero-Knowledge [NUZK] with Auxiliary Input): The simulator now depends non-uniformly on the CRS. For all CRS, there exists a circuit S_{CRS}, such that for all (x, w, Aux),

$$(x, \mathsf{CRS}, \mathsf{Prove}(\mathsf{CRS}, x, w), \mathsf{Aux}) \approx_c (x, \mathsf{CRS}, S_{\mathsf{CRS}}(x, \mathsf{Aux}), \mathsf{Aux})$$

Recall: (V)WSS allows the Prover to potentially leak out interesting information about the first message (the CRS).

This is exactly captured by the Simulator's non-uniform dependence on the CRS!

The Simulator knows something about the CRS that even the malicious Verifier does not.

Our Main Construction

Subexponential Hardness of LWE

↓

Malicious Uniform Random String (URS) NUZK Argument

↓

Two-round Public Coin (V)WSS Argument

Main Theorem (Informal): Assuming the subexponential hardness of LWE, there exists a two-round public-coin argument system that satisfies *both* WSS and VWSS.

Main Technical Tool: We construct the first ZAP with computationally adaptive soundness from the subexponential hardness of LWE.

* Requires the existence of a **Super-dense PKE** from LWE.

Super-dense PKE from LWE

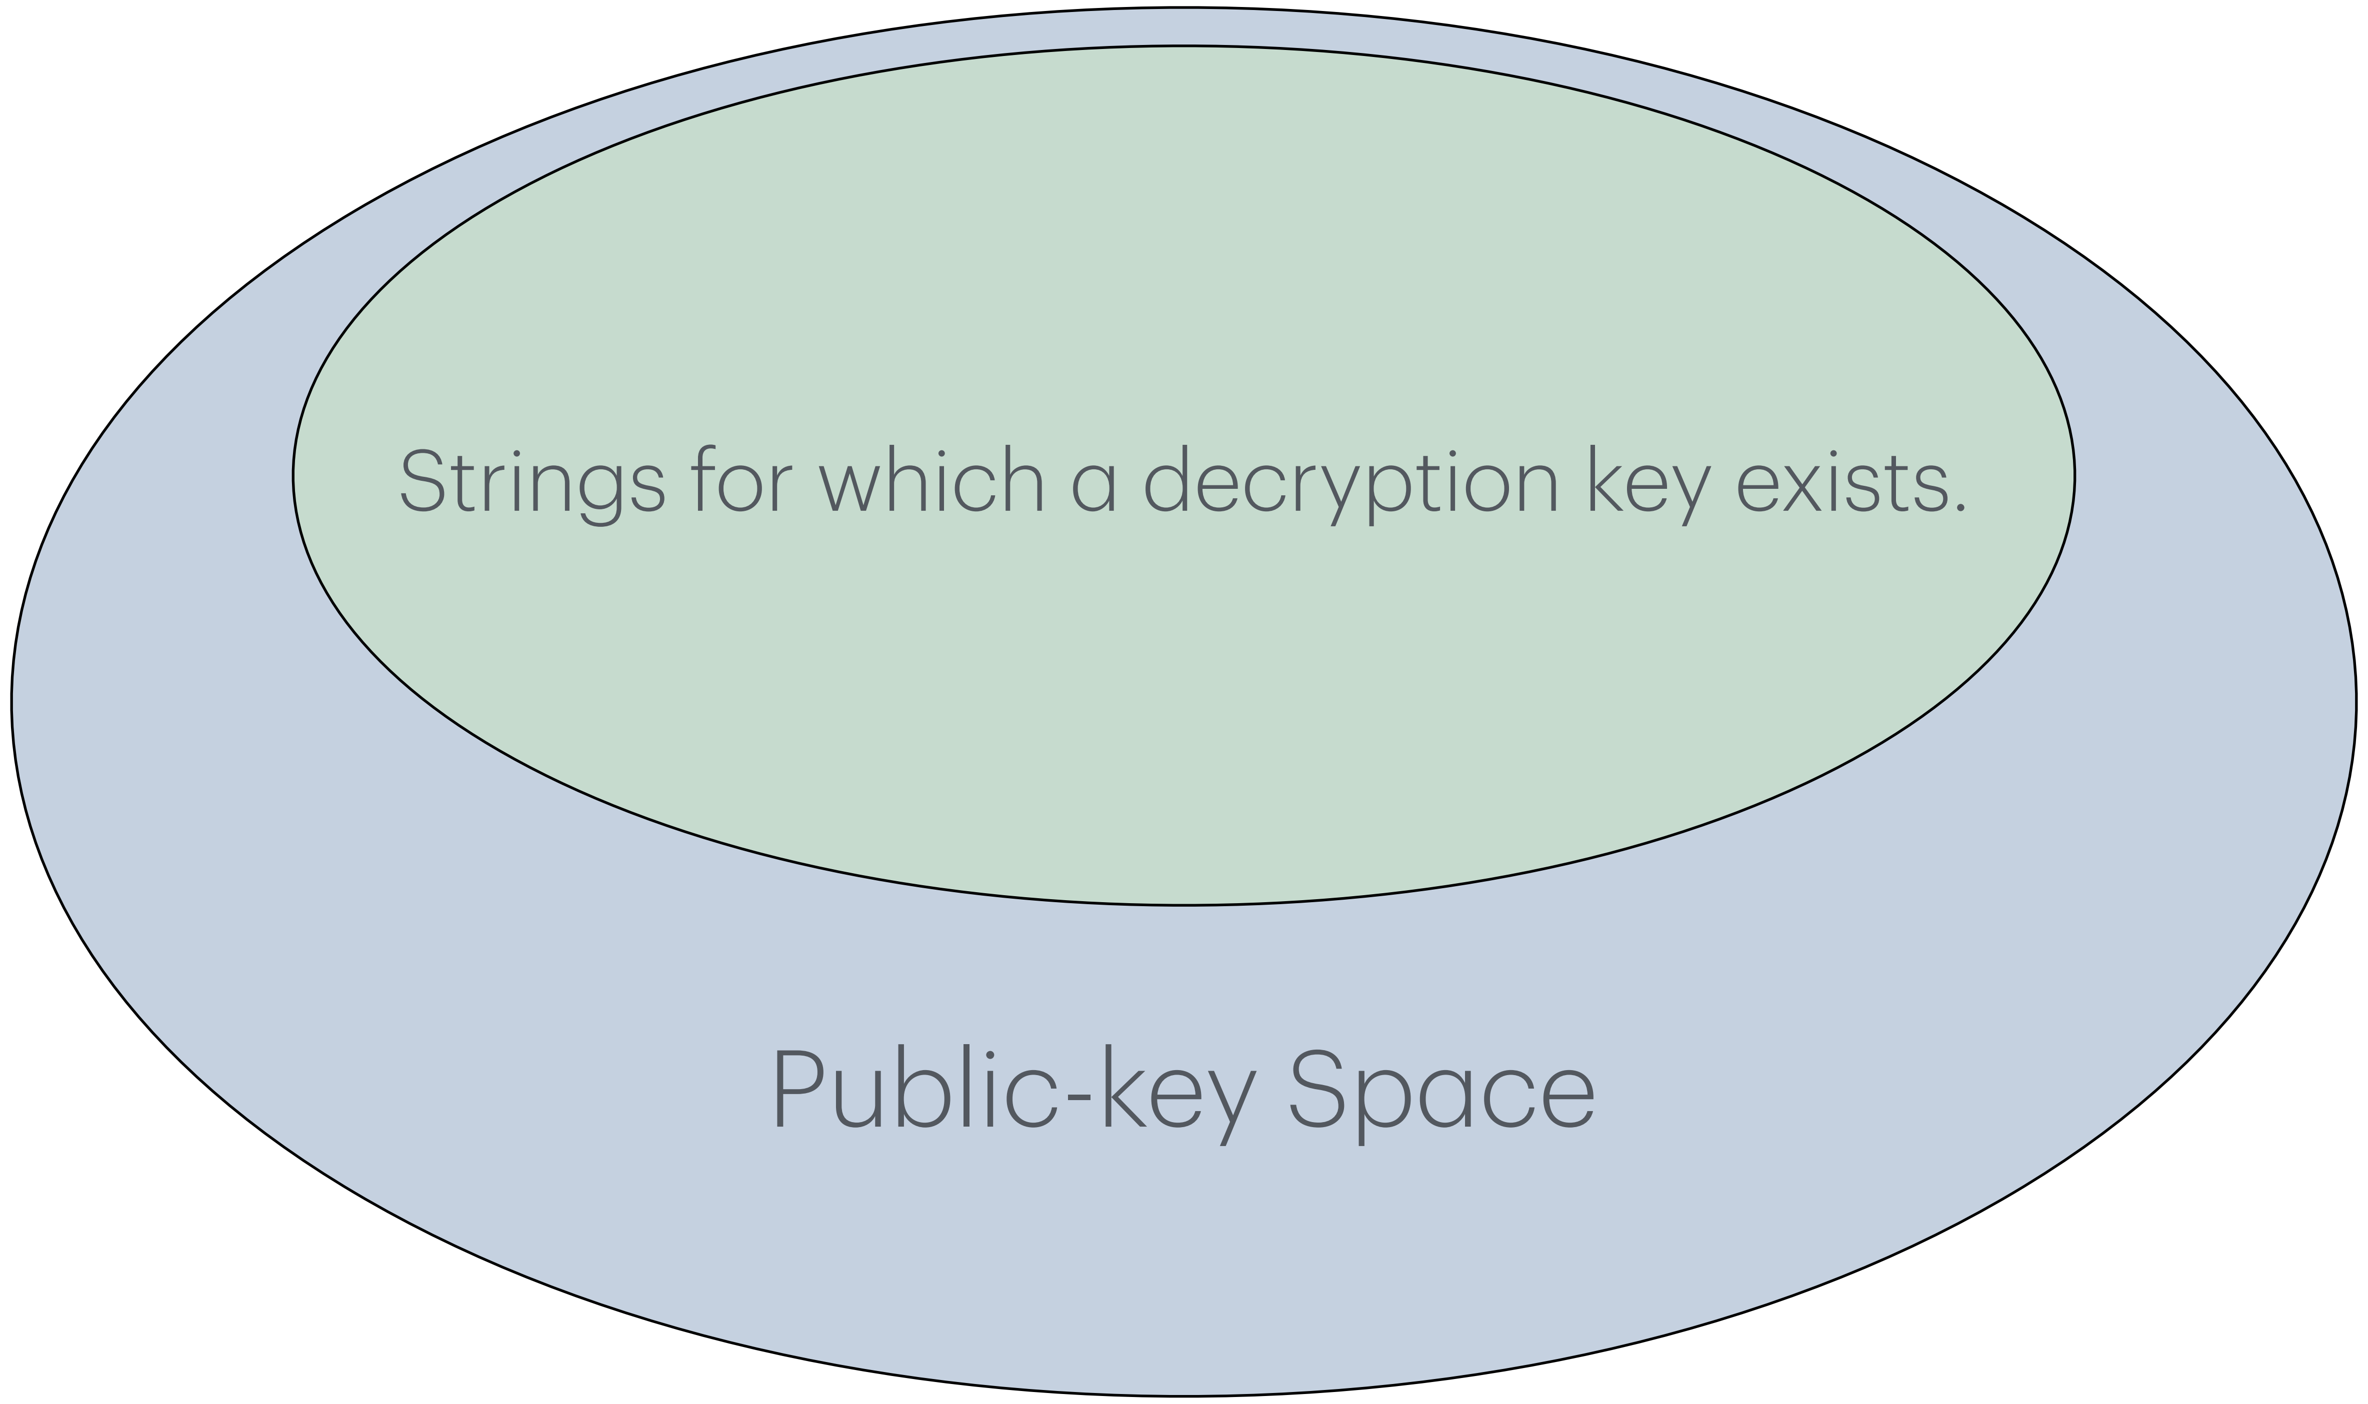

Density: The probability that a random string is a valid public key.

Super-dense PKE from LWE

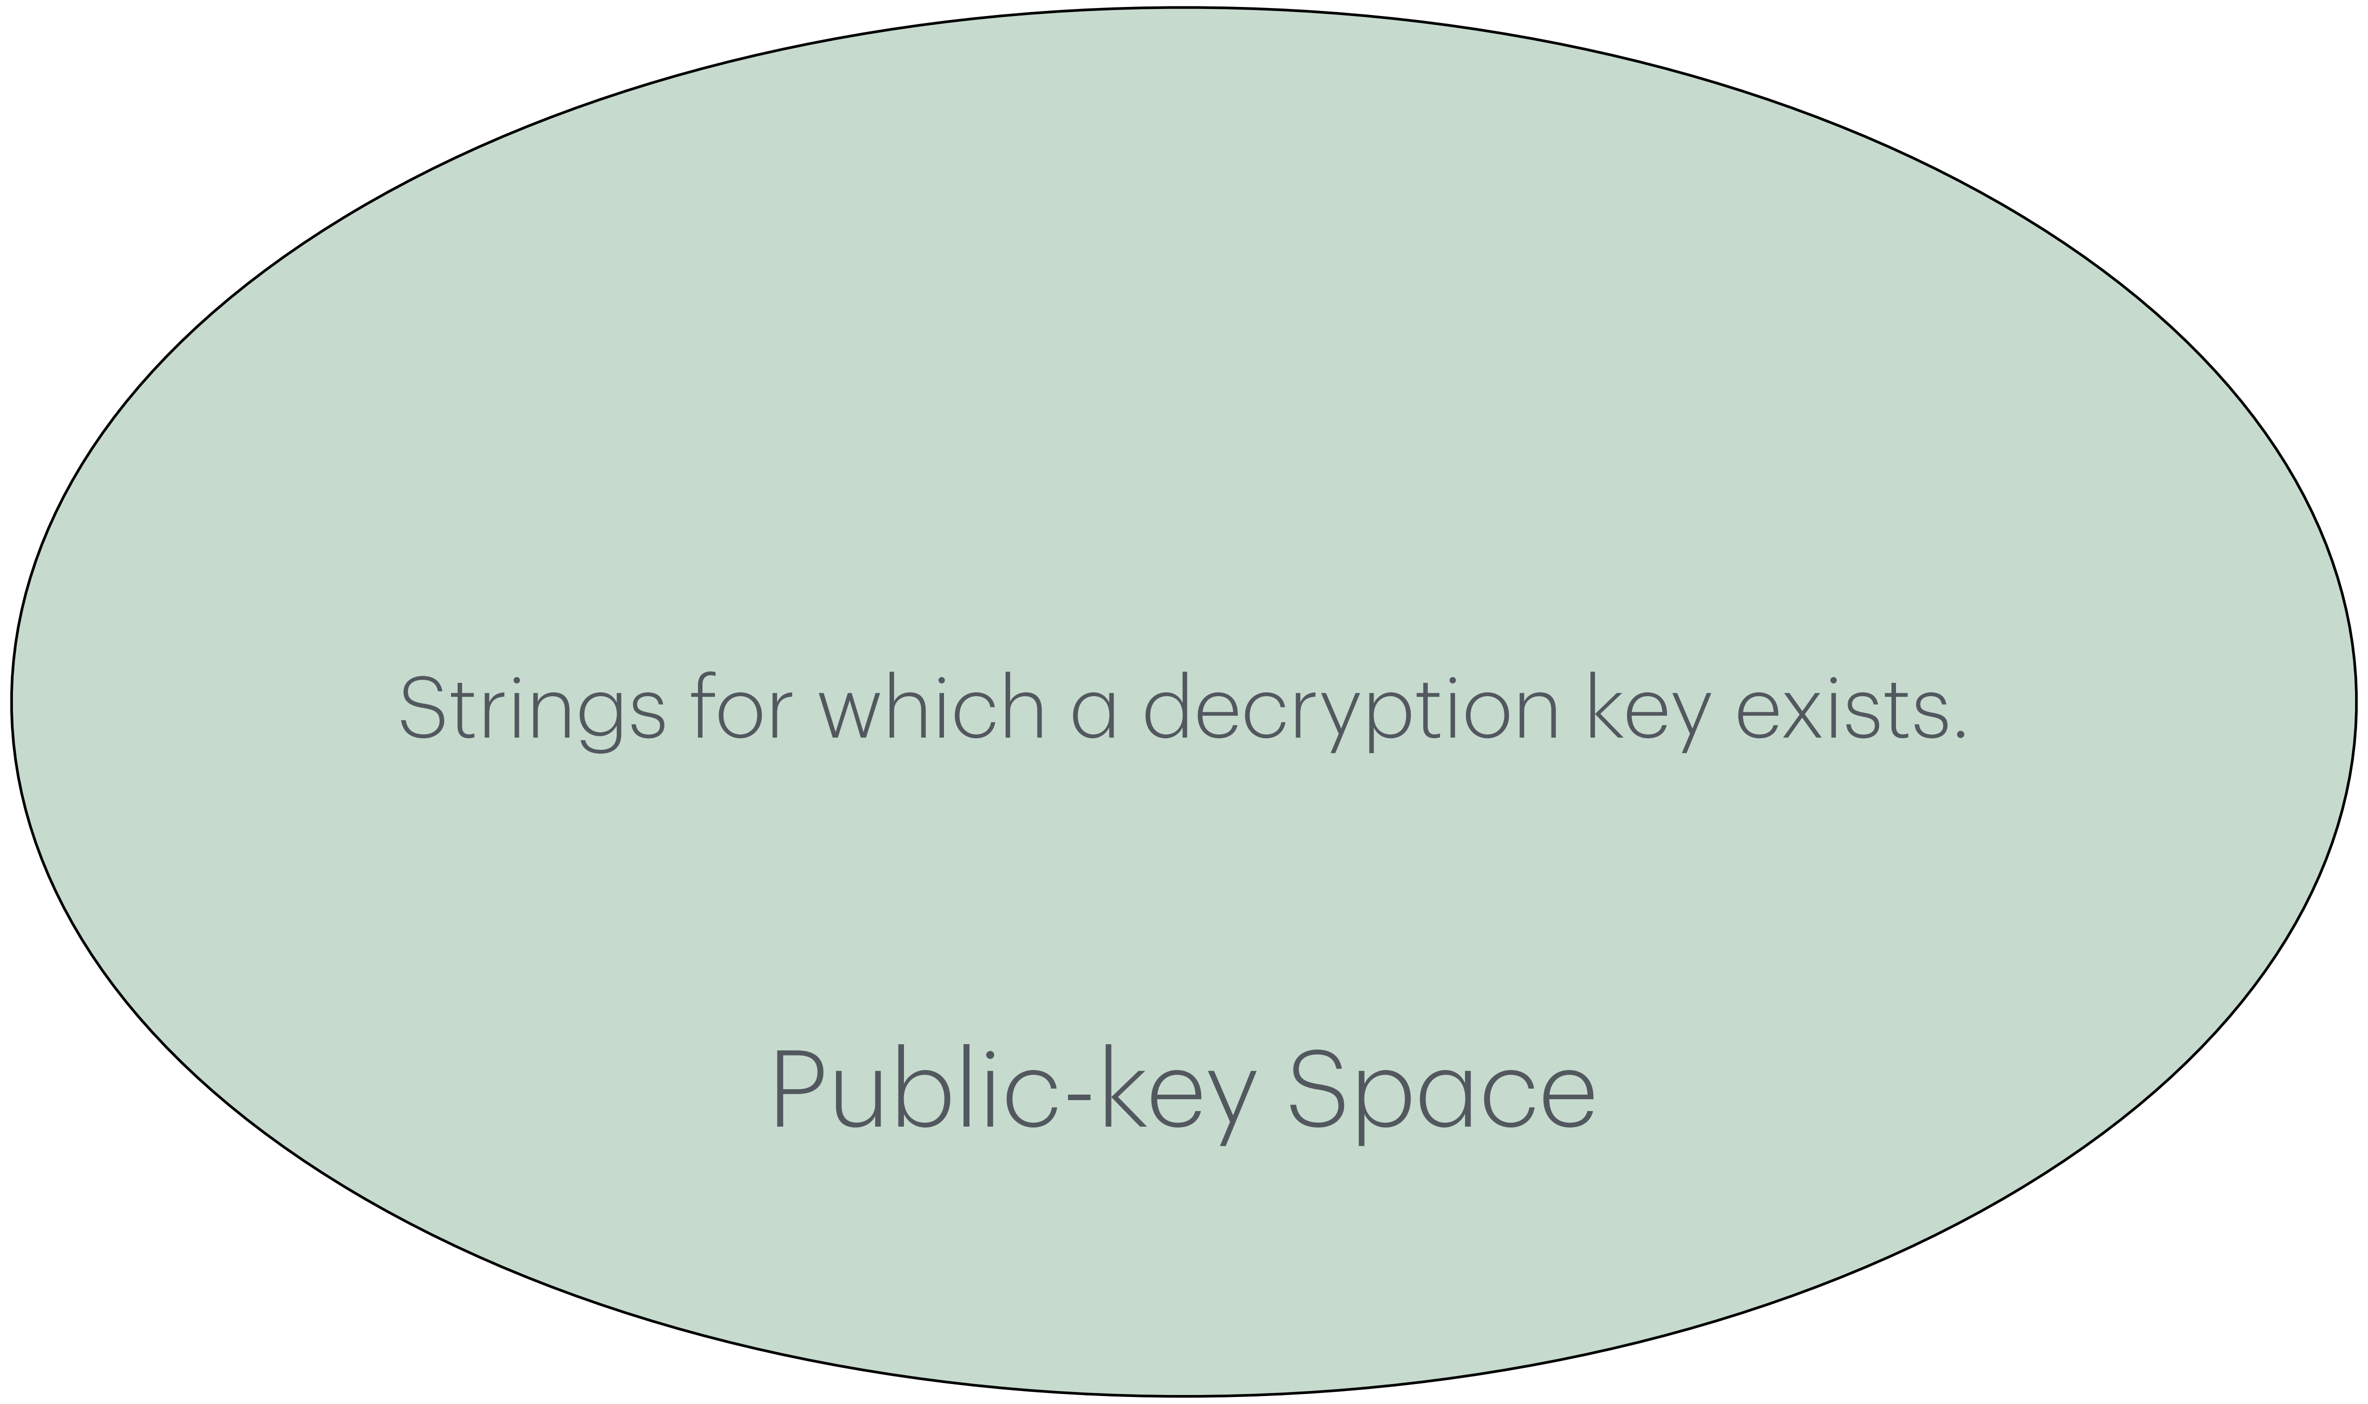

Super-dense: *All* possible strings are valid public keys.

Previously unknown from LWE (Goyal, Jain, Jin, Malavolta '20; Badrinarayan, Fernando, Jain, Khurana, Sahai '20)

Super-dense PKE from LWE

Dual Regev Encryption Scheme

Public key is of the form: $\begin{bmatrix} \mathbf{A} \\ \mathbf{r}^\top \mathbf{A} \end{bmatrix}$ where $\mathbf{r}$ is a vector of small entries over $\mathbb{F}_q$.

Decryption key: $\begin{bmatrix} \mathbf{r}^\top & -1 \end{bmatrix}$.

Encrypting a bit b: $\mathbf{ct} = \begin{bmatrix} \mathbf{A} \\ \mathbf{r}^\top \mathbf{A} \end{bmatrix} \cdot \mathbf{s} + \mathbf{e} + \begin{bmatrix} \mathbf{0} \\ b \cdot \lfloor q/2 \rfloor \end{bmatrix}$.

Super-dense PKE from LWE

Dual Regev Encryption Scheme

To decrypt, compute

$$\begin{bmatrix}\mathbf{r}^\top & -1\end{bmatrix} \cdot \left(\begin{bmatrix}\mathbf{A} \\ \mathbf{r}^\top\mathbf{A}\end{bmatrix} \cdot \mathbf{s} + \mathbf{e} + \begin{bmatrix}\mathbf{0} \\ b \cdot \lfloor q/2 \rfloor\end{bmatrix} \right)$$

...and round!

What makes a matrix a valid public key?

The existence of a short solution with a non-zero last coordinate.
Certainly not true of many matrices, so dual Regev is not super-dense.

Super-dense PKE from LWE

Our work: Super-dense Dual Regev Encryption

Modification:

Encrypting a bit b:

in the ith row

$$\left(\mathsf{ct}_i = \begin{bmatrix}\mathbf{A} \\ \mathbf{r}^\top \mathbf{A}\end{bmatrix} \cdot \mathbf{s} + \mathbf{e} + \begin{bmatrix}\mathbf{0} \\ b \cdot \lfloor q/2 \rfloor \\ \mathbf{0}\end{bmatrix}\right)_{i \in [n+1]}.$$

Super-density: For *every* $\tilde{\mathbf{A}}$, there exists some *non-zero* short solution to $\tilde{\mathbf{A}}$, which may not be of the form of the honestly generated secret keys, but allow for the same decryption guarantees.

Open Questions

- Can we obtain plain model *non-interactive* (V)WSS?
  - Related to the open standing question of plain model non-interactive witness hiding (NIWH).

Thank you!