From Succinct Arguments to Proofs: Lower Bounds for Obfuscation and SNARGs

Abhishek Jain
NTT Research & JHU

Surya Mathialagan
NTT Research

Lawrence Roy
IBM Research Zurich

Abstract

We establish barriers to basing indistinguishability obfuscation (iO) for circuits and witness encryption (WE) for NP on falsifiable assumptions via black-box security reductions, formalizing the folklore *input-length barrier*.

- **One-query reductions.** If iO for circuits or WE for NP admits a one-query, non-uniform black-box reduction to a falsifiable assumption, then $\mathsf{NP} \subseteq \mathsf{coAM}$, and the polynomial hierarchy collapses.
- **Multi-query reductions.** If iO for circuits or WE for NP admits a multi-query, uniform black-box reduction to a falsifiable assumption, then, assuming NIZKs, there exists a public key encryption scheme (and in fact WE for NP) whose security reduces to that of one-way functions *alone*.

Our results also apply to *non-adaptively* sound SNARGs for NP, a setting left open by Gentry and Wichs [STOC 2011]: a one-query reduction implies that every NP language has a coAM protocol with a subexponential-time verifier, and a multi-query uniform reduction yields public-key encryption from subexponentially secure one-way functions. Previously, only one-query reductions for SNARGs with a short CRS had been ruled out [Chen and Jin, FOCS'25].

Our starting point is the observation that all three primitives imply a two-message laconic argument. As our main technical contribution, we show how to turn such an argument with a black-box security reduction into a statistically sound (possibly multiprover) *laconic proof* for the same language. We then obtain all of our lower bounds by exploiting the strong consequences of laconic proofs.

Contents

1 Introduction **1**
- 1.1 Our Results . . . 2
- 1.2 Organization of the Paper . . . 6

2 Technical Overview **6**
- 2.1 From Reductions to Laconic Proofs . . . 7
- 2.2 From Laconic Proofs to coAM Protocols . . . 10
- 2.3 From Multi-Query Reductions to Witness Encryption . . . 11
- 2.4 Implications and Extensions . . . 13

3 Related Work **14**

4 Preliminaries **14**
- 4.1 Cryptographic Definitions and Primitives . . . 14
  - 4.1.1 Games and Security . . . 14
  - 4.1.2 Two-Message Laconic Arguments . . . 16
  - 4.1.3 Succinct Non-Interactive Arguments (SNARGs) . . . 17
  - 4.1.4 Witness Encryption . . . 18
  - 4.1.5 Indistinguishability Obfuscation . . . 20
  - 4.1.6 Black-Box Reductions . . . 20
  - 4.1.7 Other Cryptographic Tools . . . 22
- 4.2 Complexity Theory . . . 24
  - 4.2.1 Interactive Proofs and Arthur–Merlin Proofs . . . 25
  - 4.2.2 Multiprover Interactive Proofs . . . 27
- 4.3 Statistics, Hashing, and Sampling . . . 28

5 From Laconic Interactive Proofs to coAM with Polynomial Communication **30**

6 From One-Query Reductions to Laconic Interactive Proofs **37**
- 6.1 Non-Uniform Reductions for Non-Adaptive Soundness . . . 38
- 6.2 Arbitrary Advice Distributions . . . 42
- 6.3 Non-Uniform Reductions for Adaptive Soundness . . . 43

7 Implications of Black-Box Reductions **45**
- 7.1 Implications for Witness Encryption and Obfuscation . . . 45
- 7.2 Implications for SNARGs . . . 47
  - 7.2.1 Non-Adaptive Soundness . . . 47
  - 7.2.2 Adaptive Soundness . . . 49
- 7.3 Implication for Adversary-Dependent Advice . . . 49

8 From Multi-Query Reductions to Multiprover Interactive Proofs **50**
- 8.1 From One Stateless Prover to Independent Provers . . . 51
- 8.2 Construction of the MIP . . . 52

9 From MIPs to Witness and Public-Key Encryption **55**
9.1 Witness Encryption 55
9.2 Public-Key Encryption 61
9.3 Implications 61

10 AI Methodology **63**

11 References **63**

A Multi-Query Lower Bounds for Adaptive SNARGs **71**
A.1 Technical Overview 71
A.2 Tools and Definitions from Prior Work 72
A.3 From Multi-Query Reductions to Small Circuits 73
A.4 Implication for Adaptive SNARGs 75

B Comparison with Existing Constructions **76**

1 Introduction

Indistinguishability obfuscation (iO) [BGI+01, GGH+13], witness encryption (WE) [GGSW13], and succinct non-interactive arguments (SNARGs) [Kil92, Mic94] are among the most powerful primitives in cryptography. They are also among the hardest to construct. A decade of effort has yielded iO from well-founded assumptions [JLS21, JLS22, RVV24], yet all known constructions of iO and WE rest on *subexponential* hardness. Furthermore, all known constructions of SNARGs for NP rest either on iO or WE themselves [SW14, JKLM25] or on new assumptions about proof complexity [HJLM26, DHK+26]. This leaves open a basic question:

Can iO, WE, *or SNARGs for* NP *be based on standard, falsifiable assumptions?*

Following [Nao03, GW11], an assumption is *efficiently falsifiable* if its security game has a polynomial time challenger. Essentially all standard assumptions, from one-way functions to LWE, are efficiently falsifiable.

The input-length barrier. For iO and WE, a longstanding informal intuition suggests that the answer is no. The security games of iO and WE are themselves not efficiently falsifiable. To check that the adversary's challenge is valid, the challenger must decide whether two circuits are equivalent, or whether a statement lies outside an NP language; this takes time 2^m for circuits with m-bit inputs or witnesses of length m. A security reduction to a falsifiable assumption would seemingly face the same task, and hence also run in time 2^m. This intuition, often referred to as the *input-length barrier*, is consistent with all known constructions: the reductions underlying iO from well-founded assumptions [AJ15, BV15, JLS21, JLS22, RVV24] incur a loss of at least 2^m. An exciting line of research [GPSZ17, LZ17, KLMR18, JJ22, MDS25] has shown that the barrier can be circumvented for restricted classes of circuits. Whether it holds in general, however, has remained open: the barrier has never been formalized, let alone proven, and no lower bound for iO or WE against general falsifiable assumptions is known.

Lower bounds for SNARGs. For SNARGs, whose soundness is likewise not efficiently falsifiable, the intuition has been made rigorous, at least in part. In a celebrated work, Gentry and Wichs [GW11] consider *adaptively sound* SNARGs, where the cheating prover may choose its statement after seeing the common reference string (CRS). Assuming subexponentially secure one-way functions, they show that such SNARGs cannot be proven secure via a black-box reduction[1] to any efficiently falsifiable assumption. Their argument relies on adaptivity in an essential way, however, and leaves open the *non-adaptive* setting, where the statement is fixed before the CRS is sampled.

The non-adaptive setting was addressed only recently, by Chen and Jin [CJ25], who consider SNARGs whose CRS is *short*, namely sublinear in the witness length. Assuming subexponentially secure super-bit generators [Rud97], they show that such SNARGs cannot be proven secure via a one-query black-box reduction to any efficiently falsifiable assumption.[2]

Neither result yields a lower bound for iO or WE, although it is natural to expect SNARG lower bounds to do so. Indeed, the SNARGs of [SW14, JKLM25] are built from iO or WE via one-query black-box reductions,[3] so a lower bound for non-adaptively sound SNARGs against one-query reductions would translate directly into the input-length barrier. However, these SNARGs have a long CRS, as does every known SNARG with a black-box security reduction, and the result

[1] Here, "black-box" means that the reduction has oracle access to the cheating prover.

[2] The full version of [CJ25] also handles a restricted multi-query setting, in which a number of queries independent of the statement length is made as a single batched query. Our one-query results capture this setting as well (Remark 7.6), so we focus on the one-query result.

[3] More precisely, [JKLM25] rely on a WE scheme for NP with short propositional proofs of correctness.

of [CJ25] requires a short one. Whether non-adaptively sound SNARGs with a long CRS can be based on falsifiable assumptions has thus remained open, and with it the input-length barrier.

The result of [CJ25] is further restricted to reductions that query the cheating prover only once. Multi-query reductions are a natural and powerful tool: they enable knowledge extraction for UP languages, where every statement has a unique witness [CGKS23, MPV24]. Whether they could similarly be leveraged to base SNARGs, WE, or iO for all of NP on falsifiable assumptions has likewise remained open.

Why the barriers resist current techniques. Both difficulties, namely, handling a long CRS and handling multiple queries, can be traced to the one technique underlying all known lower bounds. The works of [GW11, CJ25] rely on *leakage simulation* [GW11, JP14, CLP15, CCL18]: a short proof is viewed as a small amount of leakage on the witness, and a leakage simulation lemma shows that such leakage can be efficiently simulated given only the statement, provided the statement is drawn from a sufficiently hard distribution. The simulated proofs allow an efficient algorithm to stand in for the (inefficient) cheating prover, so any reduction that breaks the assumption using the cheating prover also breaks it on its own.

The short-CRS restriction arises from how this template adapts to the non-adaptive setting. There, the reduction may choose the CRS after seeing the statement, so a simulated proof must remain indistinguishable for every such choice. Leakage simulation does not support such statement-dependent choices, so the argument instead guesses a CRS, simulates a short proof for that CRS, and succeeds if the reduction supplies the same one. This step loses a factor exponential in the CRS length, which the hardness of the language can absorb only when the CRS is short. Handling a long CRS thus requires departing from this template altogether.

The one-query restriction reflects a more basic gap. The folklore intuition is that a lower bound should follow whenever the information a reduction extracts from the cheating prover is small and independent of the witness. With multiple queries, however, the reduction obtains many proofs of the same statement under different CRSs, and ruling out that these proofs jointly reveal the witness has so far proven highly challenging.[4] Both difficulties therefore call for a new approach, which we describe next.

1.1 Our Results

In this work, we present a new approach to proving black-box lower bounds for iO, WE, and SNARGs. Our starting point is the observation that all three primitives imply a two-message *laconic argument* [BIOW20], in which the verifier sends one message and the prover replies with a message far shorter than the witness. A SNARG directly implies a laconic argument where the verifier's message consists of the CRS. Furthermore, WE, and hence iO [GGH+13], gives a laconic argument where the prover's message consists of a single bit.[5]

We prove lower bounds for laconic arguments directly, thereby obtaining lower bounds for iO, WE, and SNARGs simultaneously. Instead of simulating the cheating prover, as in the leakage-simulation template, we connect such arguments, which are *computationally sound*, to *statistically sound* proof systems. We then derive our lower bounds by exploiting the strong consequences of laconic proofs. Importantly, our results place no restriction on the length of the verifier's message.

In order to state our results, we first introduce some terminology. We model a cryptographic assumption as a game between an adversary and a challenger, who ultimately decides whether the

[4] Chen and Jin [CJ25] handle multiple queries in the related notion of incrementally verifiable computation, using a language with many witnesses and a cheating prover that uses a fresh witness on each query. This strategy, however, does not work for SNARGs.

[5] The protocol is simple: the verifier encrypts a random bit under the statement and the prover decrypts it.

adversary won; the assumption is T-*time falsifiable* if the challenger runs in time T, and efficiently falsifiable if T is polynomial (Section 4.1.1). A *black-box reduction* gets oracle access to any successful cheating prover, possibly inefficient, and uses it to break the assumption. It may query the prover but cannot inspect its code, and there is no restriction on how the construction uses the underlying primitive. We use the same bound T for the running time and the advantage of the reduction. Unless stated otherwise, reductions may use deterministic non-uniform advice depending on the statement but not on the cheating prover.[6] See Section 4.1.6 for formal definitions.

Barriers to One-Query Black-Box Reductions. We first consider *one-query* black-box reductions. Although seemingly limited, one-query reductions are used in all known constructions of SNARGs that rely on black-box reductions [SW14, CJJ21, CJJ22, BBK+23, JKLV24, JKLM25] (see [CJ25, Section 2.4]), as well as in the constructions of iO from functional encryption [AJ15, BV15].

Our first result shows that any two-message laconic argument with a one-query black-box reduction can be turned into a public-coin *proof* system for the complement of the language. Recall that $L \in \mathsf{coAM}$ if the complement of L has a two-message public-coin proof: the verifier (Arthur) sends random coins, the prover (Merlin) replies, and Arthur decides. We write $L \in \mathsf{coAM}[t]$ if, in addition, Arthur runs in time t.

Theorem 1.1 (Informal version of Corollary 6.5)**.** *Suppose the non-adaptive soundness of a two-message laconic argument for L, with prover message of length $\ell = \ell(\lambda, n)$, reduces to a T-time falsifiable assumption $\mathcal{A}$ via a T-time, one-query, non-uniform black-box reduction $\mathcal{R}$. Assuming the security of $\mathcal{A}$,*

$$L \in \mathsf{coAM}\left[2^{\ell} \cdot \mathsf{poly}(T)\right],$$

with communication and randomness $\mathsf{poly}(T)$.

In fact, Theorem 1.1 requires only the significantly weaker assumption that $\mathcal{A}$ is secure against $\mathcal{R}^{\mathcal{P}}$, that is, against the single, fixed algorithm obtained by running the reduction $\mathcal{R}$ with black-box access to the honest prover algorithm $\mathcal{P}$ of the laconic argument. This assumption is used only to establish the completeness of the resulting proof system. As observed in [RV10], such an assumption sometimes comes for free.[7]

To see why Theorem 1.1 is surprising, recall that brute force decides $\overline{L}$ in time $2^m\,\mathsf{poly}(n)$ by enumerating all witnesses of length m. The theorem says that Arthur can instead decide in time essentially 2^{ℓ} with the help of a polynomial-size certificate from Merlin. When $\ell \ll m$, this is a dramatic speedup that we do not expect to exist for hard languages. Furthermore, when ℓ is logarithmic and T is a polynomial-time function, we get that $L \in \mathsf{coAM}$.

Implications for Obfuscation, Witness Encryption and SNARGs. Applying Theorem 1.1 to the one-bit laconic arguments implied by WE and iO, we obtain the first lower bounds for these primitives against black-box reductions to arbitrary falsifiable assumptions. In particular, we formally establish the input-length barrier [GPSZ17, LZ17, JJ22] for one-query black-box reductions.

Roughly speaking, we show that if WE for NP and iO for circuits admit a one-query, black-box security reduction to efficiently falsifiable assumptions, then the polynomial hierarchy collapses [BHZ87]. More generally, we show that:

[6] Our one-query results also hold for randomized advice whose sampler has description complexity $\mathsf{poly}(T)$. Furthermore, our adaptive multi-query result also allows randomized advice.

[7] For example, suppose that $\mathcal{A}$ asserts the pseudorandomness of a function family. Further, suppose that $\mathcal{R}$ only depends on $\mathcal{A}$ in a black-box manner. Since $\mathcal{R}^{\mathcal{P}}$ is a fixed algorithm making some bounded number t of queries, the family can be replaced by a t-wise independent hash function, which is indistinguishable from random for $\mathcal{R}^{\mathcal{P}}$.

Corollary 1.2 (Informal version of Corollaries 7.2 and 7.4)**.** *Suppose* WE *for* NP *or* iO *for circuits admits a T-time, one-query, non-uniform black-box security reduction to a T-time falsifiable assumption secure against non-uniform adversaries of size* $\mathsf{poly}(T)$*. Then*

$$\mathsf{NP} \subseteq \mathsf{coAM}\big[\mathsf{poly}(T)\big],$$

with communication and randomness $\mathsf{poly}(T)$*. In particular, polynomial-time reductions to efficiently falsifiable assumptions imply* $\mathsf{NP} \subseteq \mathsf{coAM}$*, and hence a collapse of the polynomial hierarchy.*

To put this barrier in perspective, note that the best known Merlin-Arthur protocols for UNSAT [Wil16] have verifier time $2^{n/2} \cdot \mathsf{poly}(n)$ on formulas with n variables; a coAM protocol with verifier time 2^{n^ϵ} for any constant $0 < \epsilon < 1$ would be a dramatic improvement. Our result thus explains the exponential loss in input length incurred by the security reductions in known constructions of iO from functional encryption [AJ15, BV15], and shows that it is in fact inherent.[8]

We next consider the case of non-adaptive SNARGs. The prover's message in the laconic argument obtained from a SNARG has length $\mathsf{poly}(\lambda) \cdot n^{o(1)}$. Choosing $\lambda = n^{\epsilon'}$ for a sufficiently small constant ϵ' gives the following result:

Corollary 1.3 (Informal version of Corollary 7.7)**.** *Suppose a non-adaptively sound SNARG for $L \in \mathsf{NP}$ admits a T-time, one-query, non-uniform black-box security reduction to a T-time falsifiable assumption secure against non-uniform adversaries of size* $\mathsf{poly}(T)$*. Then, for every constant $\epsilon > 0$,*

$$L \in \mathsf{coAM}\big[2^{n^\epsilon} \cdot \mathsf{poly}(T)\big],$$

with communication and randomness $\mathsf{poly}(T)$*. In particular, polynomial-time reductions to efficiently falsifiable assumptions imply $L \in \mathsf{coAM}[2^{n^\epsilon}]$.*

Thus, for polynomial T, the worst-case assumption that $\mathsf{NP} \not\subseteq \mathsf{coAM}[2^{n^\epsilon}]$ for some constant $\epsilon > 0$ (see discussion above) rules out such reductions for non-adaptively sound SNARGs for NP.

Barriers to Multi-Query Black-Box Reductions. We next consider *multi-query* black-box reductions, which may query the cheating prover polynomially many times. Here, we do not obtain a complexity-theoretic consequence as in Corollaries 1.2 and 1.3, and this is to be expected: if perfectly sound unique witness maps exist, then witness encryption can be constructed with a multi-query black-box reduction to one-way functions, yet no complexity-theoretic consequences of such maps are known [CPW20, CPW23]. Witness encryption from one-way functions nevertheless seems implausible, since it would yield public-key encryption from one-way functions [GGSW13]. This suggests looking for cryptographic rather than complexity-theoretic consequences.

We obtain exactly such a consequence: a uniform multi-query reduction for iO, WE, or SNARGs would yield public-key encryption whose security reduces to one-way functions *alone*, which would resolve a long-standing open problem.

Theorem 1.4 (Informal version of Corollaries 9.8 and 9.9)**.** *Suppose* WE *for* NP *or* iO *for circuits* (*resp., non-adaptively sound SNARGs for* NP) *admits a polynomial-time, multi-query, uniform black-box reduction to a secure efficiently falsifiable assumption $\mathcal{A}$. Assume further that there exist polynomially* (*resp., subexponentially*) *secure one-way functions and perfectly complete, statistically sound NIZKs in the CRS model with polynomial* (*resp., subexponential*) *zero knowledge. Then there is a polynomial-time public-key encryption scheme whose security has a black-box reduction to the security of the one-way function alone.*

[8] Their security reductions make a single query: a hybrid argument picks a random hybrid index, embeds the functional-encryption challenge there, and runs the iO adversary once. However, the hybrids range over all 2^m inputs, where m is the input length, so these reductions lose a factor of 2^m and require subexponentially secure functional encryption. Since the security of functional encryption is efficiently falsifiable, Corollary 1.2 shows that this loss is inherent.

As in Theorem 1.1, $\mathcal{A}$ is only required to be secure against $\mathcal{R}^{\mathcal{P}}$, where $\mathcal{P}$ is the honest prover of the laconic argument; for WE and iO, this is the honest decryption algorithm. Moreover, the zero-knowledge property of the NIZK is used only for the correctness of the resulting scheme.[9] For WE and iO, the theorem extends to T-time falsifiable assumptions, and yields a scheme where the algorithms have runtime $\mathsf{poly}(T)$.

Theorem 1.4 does not contradict the black-box separations of public-key encryption from one-way functions [IR89, RTV04], since the construction uses the code of the one-way function: it appears in the statements proven using the NIZK (see Section 3). Nevertheless, since basing public-key encryption on one-way functions has resisted decades of effort, we view Theorem 1.4 as evidence that such reductions are unlikely to exist. In particular, this suggests that the input-length barrier for iO and WE holds even against multi-query, uniform black-box reductions.

In fact, we obtain the stronger implication anticipated above: under the same assumptions as Theorem 1.4, we can construct witness encryption for NP whose security reduces to one-way functions *alone*. In the case of WE and iO, the resulting witness encryption scheme runs in polynomial time. In the case of SNARGs, the prover's message is longer, and we obtain exponentially efficient witness encryption (XWE) [BJK+18]: for every constant $\epsilon > 0$, encryption and decryption take time $\mathsf{poly}(\lambda) \cdot 2^{O(n^\epsilon)}$ on statements of length n, while ciphertexts are of polynomial size. [10]

Theorem 1.4 extends to non-uniform reductions whose advice depends only on the statement length, at the cost of a non-uniform encryption algorithm in the resulting schemes (Section 9). Our results do not apply to reductions that receive statement-dependent advice: our construction invokes the reduction on statements sampled during encryption, and the encryptor has no way to obtain the advice for such statements.

Barriers to Adaptive Soundness from Worst-Case Assumptions. Finally, we consider *adaptively sound* SNARGs. Here, the barrier of Gentry and Wichs [GW11] already covers multi-query reductions, but it relies on an average-case assumption, namely subexponentially secure one-way functions. We show, with AI assistance (see Section 10), that their barrier holds under a *worst-case* assumption.

Theorem 1.5 (Informal version of Corollary 7.8)**.** *Suppose an adaptively sound SNARG for $L \in \mathsf{NP}$ admits a T-time, multi-query black-box reduction, with adaptive queries and possibly randomized advice depending only on the statement length, to a T-time falsifiable assumption secure against non-uniform adversaries of size $\mathsf{poly}(T)$. Then*

$$L \in \mathsf{SIZE}\left[2^{n^\epsilon} \cdot \mathsf{poly}(T)\right] \qquad \textit{for every constant } \epsilon > 0.$$

In particular, polynomial-time reductions to efficiently falsifiable assumptions imply $L \in \mathsf{SIZE}[2^{n^\epsilon}]$.

Thus, for polynomial T, the worst-case assumption that $\mathsf{NP} \not\subseteq \mathsf{SIZE}[2^{n^\epsilon}]$ for some constant $\epsilon > 0$ rules out such reductions for adaptively sound SNARGs for NP. Since this assumption is implied by that of Gentry and Wichs [GW11], our result strengthens theirs. Our proof adapts their argument and is given in Appendix A.

Our Approach: From Arguments to Proofs. At the heart of our results is a transformation of a two-message laconic argument with a q-query black-box reduction into a statistically sound *laconic*

[9] In fact, witness indistinguishability against a single, fixed algorithm suffices for this purpose; see Section 9.

[10] Trivial WE, in which the encryptor decides whether $x \in L$ by brute force over all 2^m witnesses of length m, requires no assumptions. Assuming subexponential LWE, [BJK+18] construct XWE with encryption and decryption time as well as ciphertext size $2^{m/2} \cdot \mathsf{poly}(\lambda, n)$. XWE with this efficiency is not known to imply public-key encryption; in contrast, the one we construct does, assuming subexponential one-way functions.

proof for the same language, with a single prover when $q = 1$ and multiple noncommunicating provers in general. The verifier of the proof plays a *mixed strategy*, either running the argument's verifier or running the reduction together with the challenger of the assumption, and we show that no prover strategy survives both tests (Section 2.1). Importantly, since the verifier samples the CRS itself, the prover's communication in the proof equals the proof length of the argument, regardless of the length of the CRS.

Our lower bounds then follow from the strong consequences of laconic proofs. For one-query reductions, we refine a theorem of Goldreich, Vadhan, and Wigderson [GVW02], which turns a laconic proof for L into a public-coin proof for $\overline{L}$ with complexity exponential only in the prover's communication. Our refinement keeps the communication and randomness polynomial, which gives Theorem 1.1. For multi-query reductions, we build witness encryption from our multiprover proof. The construction relies on a *strong local correctness* property of the proof, which guarantees completeness even when each prover holds its own witness, and on a canonicalization of witnesses via NIZK proofs of an OR statement, which lets the encryptor and the decryptor run the provers on indistinguishable witnesses. It builds upon, but differs significantly from, the constructions of encryption from laconic zero-knowledge arguments [BDRV18, LMP25], where the zero-knowledge simulator plays the role of local correctness. We give a detailed overview in Section 2.

Takeaways and Open Problems. Our results suggest a simple takeaway for constructing iO, WE, or SNARGs for NP from falsifiable assumptions: neither multiple queries to the adversary nor, for SNARGs, a long CRS helps. Instead, one must either rely on subexponential hardness, with a reduction that can decide circuit equivalence (for iO) or the language (for WE and SNARGs) by brute force, or use the adversary in a non-black-box way. All known constructions take one of these two routes; see Appendix B.

Our results also leave two gaps. First, for one-query reductions, Theorem 1.1 requires deterministic advice. For randomized advice (and advice that may depend on the cheating prover), we obtain the weaker bound $L \in \mathsf{coAM}[2^{O(\ell+|\mathsf{crs}|)} \cdot \mathsf{poly}(T)]$ (Corollaries 6.7 and 7.10), which is meaningful only when the verifier's message is short. This recovers the lower bound of [CJ25] under a worst-case assumption implied by theirs. Second, for multi-query reductions, Theorem 1.4 does not cover reductions whose advice depends on the statement. Whether such reductions can also be ruled out, or instead exploited, remains open.

AI Methodology. See Section 10.

1.2 Organization of the Paper

In Section 2, we give an overview of our techniques. We discuss related work in Section 3. We give preliminaries in Section 4. In Section 5, we show how to turn laconic proofs into coAM protocols with polynomial communication. In Section 6, we show how one-query reductions yield laconic proofs, for both non-adaptive and adaptive soundness, and in Section 7, we derive the implications for witness encryption, obfuscation, and SNARGs. In Section 8, we show how multi-query reductions yield multiprover interactive proofs with strong local correctness, and in Section 9, we use them to construct witness encryption and public-key encryption. We describe our use of AI in Section 10. Appendix A contains our multi-query lower bound for adaptively sound SNARGs.

2 Technical Overview

We consider arbitrary two-message laconic arguments $\Pi = (\mathcal{V} = (\mathcal{V}_1, \mathcal{V}_2), \mathcal{P})$ for a language $L \in \mathsf{NP}$. The verifier sends a first message pk, keeping a private state vk; the prover replies with a short

message π of $\ell_{\mathcal{P}}$ bits; and the verifier decides using vk. The first message may be arbitrarily long. We allow either of the following soundness definitions (Definition 4.6):

- *Search security:* no efficient prover is accepted on a false statement with nonnegligible probability.
- *Indistinguishability security:* no efficient prover is accepted on a false statement with probability noticeably greater than $1/2$.

This template captures all three of our targets. A non-adaptively sound SNARG is a laconic argument whose first message is the CRS. Witness encryption gives one with a *one-bit* reply: the verifier encrypts a random bit under x, and the prover decrypts it, so on a false statement the soundness game *is* the security game of witness encryption, and every security reduction carries over with the same number of queries. Via [GGH+13], the same holds for iO (Section 7.1).

Our main idea is to turn a black-box security reduction for Π into a statistically sound *laconic proof*, in which the prover still sends only $\ell_{\mathcal{P}}$ bits (Section 2.1). Laconic proofs are surprisingly powerful. For one-query reductions, we complement the proof into a coAM protocol whose verifier runs in time $2^{\ell_{\mathcal{P}}} \operatorname{poly}(\lambda)$ (Section 2.2). For witness encryption and iO, a single bit means this is polynomial, and we get $\mathsf{NP} \subseteq \mathsf{coAM}$; for SNARGs, we get $\mathsf{NP} \subseteq \mathsf{coAM}[2^{n^\epsilon}]$. For multi-query reductions, we obtain a multiprover proof, and use it to build witness encryption and public-key encryption from one-way functions (Section 2.3).

2.1 From Reductions to Laconic Proofs

As a warm-up, we first consider a uniform, one-query, polynomial-time black-box reduction $\mathcal{R}$ from the soundness of Π to a cryptographic assumption $\mathcal{A}$. We assume that Π has indistinguishability security, which is the weaker of the two notions, and, for simplicity, that $\mathcal{A}$ is also an indistinguishability assumption: no efficient adversary makes $\mathcal{A}$ output win with probability noticeably greater than $1/2$.

Our goal is to show that L then has a *laconic proof*: a two-message, private-coin proof system Π' in which the prover sends the same $\ell_{\mathcal{P}}$ bits as in Π, but whose soundness is *statistical*. The verifier's message may still be long.

On a false statement x, the reduction $\mathcal{R}^{\mathcal{P}^*}(1^\lambda, x)$ interacts with $\mathcal{A}$, makes a single query pk to a cheating prover $\mathcal{P}^*$, and receives $\pi = \mathcal{P}^*(\mathsf{pk})$. A cheating prover breaks soundness if it is accepted with probability noticeably greater than $1/2$ on an honestly generated pk. Fixing a constant margin, the reduction guarantees that if $\mathcal{P}^*$ is accepted with probability at least $3/4$, then $\mathcal{R}^{\mathcal{P}^*}$ wins with probability at least $1/2 + \delta$, for some inverse polynomial δ. (That a single such δ exists for all false statements and all cheating provers is a subtle point, which we return to below.) Since $\mathcal{R}$ is black-box, this guarantee holds for *every* cheating prover, including inefficient ones.

Accepting proofs should help the reduction only on false statements. Morally, accepting proofs from the honest prover should not help the reduction, while accepting proofs from a cheating prover should. Indeed, suppose that $\mathcal{A}$ is secure, and consider the honest prover on a true statement $x \in L$. Whether or not the reduction's query looks like an honestly generated pk, the honest prover is efficient given the witness w. Thus $\mathcal{R}$, with its query answered by $\mathcal{P}(\mathsf{pk}, x, w)$, is an efficient non-uniform adversary against $\mathcal{A}$ with x and w hardwired. We denote this adversary by $\mathcal{R}^{\mathcal{P}}$; all we need from $\mathcal{A}$ is that $\mathcal{R}^{\mathcal{P}}$ does not break it. On a false statement $x \notin L$, on the other hand, the reduction guarantee applies to every prover whose proofs are accepted with probability $\alpha \geq 3/4$.

We therefore have

$$\begin{aligned} x \in L: &\quad \Pr\big[\mathcal{R}^{\mathcal{P}(\cdot,x,w)}(1^\lambda, x) \text{ wins}\big] \le \tfrac{1}{2} + \mathsf{negl}(\lambda) \quad (\mathcal{R}^{\mathcal{P}} \text{ does not break } \mathcal{A}), \\ x \notin L,\ \alpha \ge \tfrac{3}{4}: &\quad \Pr\big[\mathcal{R}^{\mathcal{P}^*}(1^\lambda, x) \text{ wins}\big] \ge \tfrac{1}{2} + \delta \qquad\qquad (\text{reduction guarantee}). \end{aligned}$$

The ideal test. This suggests the following test: accept if the proof is accepting *and* the reduction does not break the assumption. Suppose for a moment that Π is publicly verifiable, so that the verifier can check a proof π against any pk. The verifier then simulates $\mathcal{R}(1^\lambda, x)$ interacting with $\mathcal{A}$, forwards the query pk to the prover, and returns the reply π to $\mathcal{R}$. It accepts if π is an accepting proof for pk and $\mathcal{A}$ does not output win. One would hope that the honest prover passes both parts on $x \in L$, and that on $x \notin L$, a cheating prover either fails the first part or, by the second line above, helps the reduction win. Moreover, since the reduction guarantee holds against unbounded provers, this argument would give *statistical* soundness.

However, the reduction guarantee measures α on *honestly generated* pk, while the test checks the proof on the reduction's query, which may be distributed arbitrarily. A cheating prover can return accepting proofs on the reduction's queries and non-accepting proofs on honestly generated pk. Then $\alpha < 3/4$, the reduction guarantee does not apply, and $\mathcal{R}$ need not win, so the test accepts. Moreover, the arguments we care about are not publicly verifiable: for witness encryption, checking a reply requires the encrypted bit, and the reduction's query comes with no such bit at all.

The laconic proof. We therefore separate the two parts, and instead have the verifier send a message drawn from a *mixture* of two distributions: an honestly generated pk, which it uses to check the proof, and the reduction's query, which it uses to run the reduction. The prover sees only pk, so its behavior in both tests is described by a single response strategy $\mathsf{pk} \mapsto \pi$. The reduction guarantee applies to this strategy, even if it responds differently to messages from the two distributions. This gives the following laconic proof $\Pi' = (\mathcal{P}', \mathcal{V}')$ for L, in which the honest prover is unchanged.

On common input x, with the prover holding a witness w:

1. *Verifier's message.* $\mathcal{V}'$ samples a private bit b with $\Pr[b = 0] = 4/5$.
   - If $b = 0$, it samples $(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x)$.
   - If $b = 1$, it simulates $\mathcal{R}(1^\lambda, x)$ interacting with $\mathcal{A}$ until $\mathcal{R}$ makes its query pk, and saves the state of the simulation.

   It sends pk, keeping b and its state private.
2. *Prover's message.* $\mathcal{P}'$ replies with $\pi \leftarrow \mathcal{P}(\mathsf{pk}, x, w)$, of length $\ell_{\mathcal{P}}$.
3. *Verifier's decision.*
   - If $b = 0$, $\mathcal{V}'$ outputs $\mathcal{V}_2(\mathsf{vk}, x, \pi)$.
   - If $b = 1$, $\mathcal{V}'$ resumes the simulation with π as the oracle answer, and accepts if and only if $\mathcal{A}$ does not output win.

Completeness. Fix $x \in L$. When $b = 0$, the honest prover is accepted with probability $1 - \mathsf{negl}(\lambda)$ by completeness of Π. When $b = 1$, the first line of the display above shows that $\mathcal{A}$ outputs win with probability at most $1/2 + \mathsf{negl}(\lambda)$. Thus $\mathcal{V}'$ accepts with probability at least

$$\tfrac{4}{5}\big(1 - \mathsf{negl}(\lambda)\big) + \tfrac{1}{5}\big(\tfrac{1}{2} - \mathsf{negl}(\lambda)\big) = \tfrac{9}{10} - \mathsf{negl}(\lambda).$$

Soundness. Fix $x \notin L$ and an arbitrary, possibly unbounded prover $\mathcal{P}^*$. Let α be the probability that $\mathcal{V}$ accepts $\mathcal{P}^*(\mathsf{pk})$ on an honestly generated pk.

We consider two cases, according to how $\mathcal{P}^*$ responds to honestly generated pk.

(a) $\alpha < 3/4$. Here the reduction guarantee does not apply, and $\mathcal{R}$ need not win. However, $\mathcal{V}'$ catches $\mathcal{P}^*$ often enough when $b = 0$, and accepts with probability at most

$$\tfrac{4}{5} \cdot \tfrac{3}{4} + \tfrac{1}{5} = \tfrac{4}{5}.$$

(b) $\alpha \geq 3/4$. Here $\mathcal{P}^*$ may pass when $b = 0$. However, the reduction guarantee now applies, so $\mathcal{A}$ outputs win with probability at least $1/2 + \delta$ when $b = 1$. Thus $\mathcal{V}'$ accepts with probability at most

$$\tfrac{4}{5} + \tfrac{1}{5}\big(\tfrac{1}{2} - \delta\big) = \tfrac{9}{10} - \tfrac{\delta}{5}.$$

In either case, one of the two parts of the ideal test catches $\mathcal{P}^*$, and the acceptance probability is smaller than $9/10$ by an inverse polynomial.

A subtlety: a common advantage bound. Case (b) uses a single inverse polynomial δ for *all* false statements and *all* cheating provers. The black-box guarantee does not directly give this, since the polynomial it provides may depend on the family of cheating provers. We obtain a common bound by a diagonalization argument in the spirit of [Bel02] (Lemma 6.2). If none existed, we could stitch provers with ever smaller advantage into a single family that breaks Π, using a dummy prover that returns a fixed response at all other parameters. The reduction would then have negligible advantage against this family, since with the dummy prover it is again an instance of $\mathcal{R}^{\mathcal{P}}$. The stitched provers may be unbounded, so δ holds for all cheating provers.

Thus L has a two-message, private-coin, statistically sound proof with an inverse-polynomial gap, a polynomial-time verifier, and prover communication $\ell_{\mathcal{P}}$ (Theorem 6.3). The transformation incurs no exponential cost in $\ell_{\mathcal{V}}$; for witness encryption and iO, the prover still sends a single bit.

Non-uniform reductions. We now allow $\mathcal{R}$ to receive deterministic advice $z = F(1^\lambda, x)$ of polynomial length, which may depend on x but not on the cheating prover, and need not be computable. Many of the reductions in [JKLV24, JKLM25] are of this form. Since $\mathcal{V}'$ cannot compute z, we give it as part of the input, and obtain a proof only for the promise problem

$$L^F_{\text{yes}} = \big\{(1^\lambda, x, z) : x \in L\big\}, \qquad L^F_{\text{no}} = \big\{(1^\lambda, x, F(1^\lambda, x)) : x \notin L\big\}.$$

Completeness holds for *all* z, since $\mathcal{R}^{\mathcal{P}}$ with any z hardwired does not break $\mathcal{A}$, while soundness holds for the correct advice exactly as before. This guarantee for all z will be essential for our implication to coAM.

Multi-query reductions. We now allow $\mathcal{R}$ to make k queries, possibly adaptively, and send each query to a different noncommunicating prover. However, a prover's reply may depend on its *identity*, so the provers together need not behave like the single stateless oracle that the reduction guarantee quantifies over. We fix this by sending each query to a uniformly random unused prover from a pool of $k' = \Theta(k^2/\delta)$ provers (Lemma 8.2). This is indistinguishable from a single oracle that answers each query with a random prover, unless two of the k identities collide, which happens with probability at most $\binom{k}{2}/k'$. The mixed verifier is then as before, with $b = 0$ sending an honestly generated pk to a random prover, and the analysis is unchanged, with α now the average acceptance probability of the provers (Theorem 8.3). Each contacted prover still sends $\ell_{\mathcal{P}}$ bits, so for witness encryption and iO, each sends a single bit (Corollary 8.5). In fact, completeness holds even when the honest provers use *different* witnesses, since the reduction with all of their witnesses hardwired is still an instance of $\mathcal{R}^{\mathcal{P}}$. This *strong local correctness* (Definition 8.1) is what we use to build witness encryption in Section 2.3.

2.2 From Laconic Proofs to coAM Protocols

We now turn the laconic proof from Section 2.1 into a public-coin proof for the complement of L. We first recall the connection established by Goldreich, Vadhan, and Wigderson [GVW02]. They show that if L has an interactive proof with completeness $2/3$ and soundness $1/3$ in which the prover sends at most $\ell_{\mathcal{P}}$ bits in total, then $\overline{L}$ has a constant-round public-coin proof whose complexity depends only exponentially on $\ell_{\mathcal{P}}$. In particular, if the starting proof has two messages, then $\overline{L}$ has such a proof with verifier time and communication $2^{O(\ell_{\mathcal{P}})}\,\mathsf{poly}(n)$.

This does not quite work for us, even when $\ell_{\mathcal{P}} = 1$. Our laconic proof only has an inverse-polynomial completeness–soundness gap Δ, and the theorem of [GVW02] requires constant error. Amplifying the gap by parallel repetition takes about $1/\Delta^2$ copies, which multiplies the prover's message length, and hence the exponent. For witness encryption and iO, the prover would then send $\mathsf{poly}(\lambda)$ bits instead of one, and the resulting verifier would run in exponential time. In addition, for SNARGs, we need polynomial communication, with only the verifier's running time growing as $2^{\ell_{\mathcal{P}}}$.

We therefore refine the theorem of [GVW02] for two-message proofs with arbitrary completeness and soundness errors ϵ_c and ϵ_s (Theorem 5.1). Writing $\Delta = 1-\epsilon_c-\epsilon_s$, the resulting public-coin proof for $\overline{L}$ has verifier time $2^{\ell_{\mathcal{P}}}\,\mathsf{poly}(\lambda, 1/\Delta)$, and communication and randomness $\mathsf{poly}(\lambda, 1/\Delta)$. The first difference from [GVW02] is that we amplify only at the end: we complement the laconic proof with its inverse-polynomial gap as is, and amplify the resulting protocol for $\overline{L}$, where the prover's messages need not be short. For $\ell_{\mathcal{P}} = 1$, the verifier time is therefore $2 \cdot \mathsf{poly}(\lambda, 1/\Delta)$, which is polynomial.

The idea. Our transformation follows [GVW02] closely. For each verifier message pk, let S_{pk} be the set of verifier tapes ρ with $\mathcal{V}_1(x;\rho) = \mathsf{pk}$, and let

$$q(\mathsf{pk}) = \min_{\pi \in \{0,1\}^{\ell_{\mathcal{P}}}} \Pr_{\rho \leftarrow S_{\mathsf{pk}}} \left[\mathcal{V}_2(\rho, x, \pi) = 0\right]$$

be the rejection probability of the best reply to pk. Then $\mathbb{E}_{\mathsf{pk}}[q(\mathsf{pk})] \geq 1-\epsilon_s$ if $x \notin L$, and $\mathbb{E}_{\mathsf{pk}}[q(\mathsf{pk})] \leq \epsilon_c$ if $x \in L$. We describe a protocol for $\overline{L}$ between Arthur and Merlin, assuming for now that Arthur's coins are private; we later make it public-coin using generic transformations. Given approximately uniform samples from S_{pk} for many honestly generated pk, Arthur could estimate this expectation himself, by brute force over all $2^{\ell_{\mathcal{P}}}$ replies. This is the only step that takes time $2^{\ell_{\mathcal{P}}}$.

Sampling with Merlin's help. Arthur cannot sample from S_{pk} on his own, so, as in [GVW02], Merlin helps him using hashing. For each of N honestly generated messages $\mathsf{pk}_i = \mathcal{V}_1(x;\rho_i)$, Arthur sends a hash function h_i, with range set by Merlin's claimed size of S_{pk_i}, together with $h_i(\rho_i)$. Merlin replies with the bucket of tapes in S_{pk_i} with this hash, of size at most B. Since the bucket must contain the hidden tape ρ_i, a bucket with missing tapes is caught with probability $\Omega(1/B)$, and a full bucket gives an unbiased estimate of the rejection probability of each *fixed* reply, even if Merlin lies about sizes. The second difference from [GVW02] is that Arthur uses each bucket to evaluate all $2^{\ell_{\mathcal{P}}}$ replies at once, so the replies cost him running time but not communication. Soundness of the protocol for $\overline{L}$ needs only one reply per pk_i: for $x \in L$, Arthur's minimum is at most his estimate for the reply attaining $q(\mathsf{pk}_i)$. Completeness, however, needs the estimates for all replies to be accurate simultaneously. We obtain this by a union bound: with $\Theta(\ell_{\mathcal{P}} + \log(1/\Delta))$-wise independent hashing, buckets of size $B = \mathsf{poly}(\ell_{\mathcal{P}}, 1/\Delta)$ suffice. In other words, exponentially many replies cost only a factor of $\ell_{\mathcal{P}}$ in the bucket size, and $N = \mathsf{poly}(1/\Delta, \log B)$ samples suffice.

The resulting private-coin protocol for $\overline{L}$ has gap $\Omega(1/B)$, which is inverse polynomial. Using well-known results on amplifying the completeness–soundness gap by parallel repetition, trans-

forming private-coin proofs into public-coin ones [GS86], and collapsing the rounds of Arthur–Merlin proofs [Bab85], we obtain the claimed two-message public-coin proof for $\overline{L}$.

Dealing with non-uniform advice. For non-uniform reductions, we only have a laconic proof for the promise problem $(L^F_{\text{yes}}, L^F_{\text{no}})$, and the above transformation gives a public-coin proof for its complement on input $(1^\lambda, x, z)$. To obtain a proof for $\overline{L}$, we let Merlin first send the advice z. If $x \notin L$, the honest Merlin sends the correct advice $F(1^\lambda, x)$, and Arthur accepts. If $x \in L$, then $(1^\lambda, x, z) \in L^F_{\text{yes}}$ for *every* z, since completeness of the laconic proof holds for all z, so Arthur rejects whatever z Merlin sends. Collapsing rounds again [Bab85] gives a two-message public-coin proof for $\overline{L}$. In other words, completeness for all z becomes soundness of the proof for $\overline{L}$. Randomized advice $z \leftarrow \mathcal{D}_{\lambda,x}$ is handled in the same way: Merlin sends a description of a sampler for $\mathcal{D}_{\lambda,x}$, and completeness holds for every description. Advice that depends on the cheating prover reduces to this case. Both the reduction's winning probability and the acceptance probability are linear in the prover's strategy, so by the minimax theorem [Sio58] a single distribution works against all cheating provers, and by Carathéodory's theorem [COA20, Theorem 2] it has a sampler of size $2^{\ell_{\mathcal{P}}+\ell_{\mathcal{V}}}\,\mathsf{poly}(\lambda)$. The resulting public-coin proof has complexity $2^{O(\ell_{\mathcal{P}}+\ell_{\mathcal{V}})}\,\mathsf{poly}(\lambda)$, which now depends on the CRS length; in the short-CRS regime, it recovers the one-query lower bound of [CJ25] from a worst-case hypothesis (Corollaries 6.7 and 7.10).

Comparison with leakage simulation. In spirit, the Arthur–Merlin protocol above plays an analogous role to the leakage simulation lemma used in [GW11, CJ25]: both exploit that a short proof can be brute-forced at cost $2^{\ell_{\mathcal{P}}}$. The difference lies in what is required. The leakage simulation lemma produces simulated proofs that are indistinguishable from real ones to *every* efficient test, and on *every* CRS that the reduction might query. Arthur only needs to handle a single test, namely the verifier $\mathcal{V}'$, which runs the reduction against the challenger, and only on CRSs that he samples himself. Intuitively, this is why our bounds have no exponential dependence on the CRS length.

2.3 From Multi-Query Reductions to Witness Encryption

For multi-query reductions, Section 2.1 gives a multiprover proof, and the complementation step above no longer applies. Instead, we use the multiprover proof to build witness encryption. Recall that its provers are contacted at most once each, that each contacted prover sends $\ell_{\mathcal{P}}$ bits, and that the proof has strong local correctness. With the security parameter chosen as in Section 2.4, we have $\ell_{\mathcal{P}} = 1$ for witness encryption and iO (Corollary 8.5), and $\ell_{\mathcal{P}} = n^\epsilon$ for SNARGs, for an arbitrarily small constant $\epsilon > 0$ (Corollary 8.4).

Prior work: witness encryption from laconic zero knowledge. We follow [BDRV18, LMP25], who build witness encryption from laconic zero-knowledge arguments. It suffices to build a *weak* witness encryption scheme: the encryptor samples a string μ and a ciphertext, the decryptor recovers μ with high probability, and on a false statement μ is somewhat hard to predict from the ciphertext. Such a scheme can be amplified to full witness encryption [LMP25, Proposition 4.2]. In [LMP25], the encryptor uses the zero-knowledge simulator to sample a verifier message and the prover's reply distribution, while the decryptor uses its witness to sample the same reply distribution. Since replies are short, both estimate this distribution by a histogram and agree on μ by correlated sampling, using a shared random index v included in the ciphertext.

Replacing the simulator. Our proof system has no simulator: the encryptor holds no witness for x, so it cannot run the honest provers at all. We replace the simulator by a trapdoor statement, using the standard OR trick of Feige, Lapidot, and Shamir [FLS90]. To encrypt under x, the encryptor samples a seed s, sets $y = \mathsf{PRG}(s)$, and samples a CRS crs for a statistically sound NIWI. It then

considers the statement $X = (\mathsf{crs}, x, y)$ of the relation $\mathcal{R}_{\mathsf{NIWI}}$, whose witnesses are valid NIWI proofs, under crs, of the statement "$x \in L$ or y is in the image of PRG." The encryptor can produce such NIWI proofs using s, and the decryptor can produce them using its witness w for x. We now apply the multiprover proof to $\mathcal{R}_{\mathsf{NIWI}}$.

- The encryptor runs the verifier on X, answering each question with the honest prover holding a freshly sampled NIWI proof under s. It picks a uniformly random contacted prover i, with its question u. It then estimates the distribution of the honest reply of prover i to u by a histogram of many samples, each using a fresh NIWI proof, and uses correlated sampling with a random index v to select a reply μ from this histogram. The ciphertext is $(x, \mathsf{crs}, y, i, u, v)$.
- The decryptor estimates the same reply distribution, using NIWI proofs computed from w instead of s, and applies correlated sampling with the same index v to its own histogram.

Security. Suppose $x \notin L$, and that an adversary $\mathcal{B}$ predicts μ from the ciphertext. We define cheating provers from $\mathcal{B}$: on question u, prover i samples v and outputs $\mathcal{B}(x, \mathsf{crs}, y, i, u, v)$. Consider running the verifier with these provers in two experiments.

- If $y = \mathsf{PRG}(s)$, consider the encryptor's execution of the verifier, in which each prover holds its own, independently sampled NIWI proof. By strong local correctness, which allows every prover to hold a different witness, this execution is accepted with probability at least $1 - \epsilon_c$. The ciphertext describes a uniformly random contacted prover i with its question u, and μ is a sample of its reply. Thus, if $\mathcal{B}$ errs with probability η, a coupling argument that replaces the honest replies by the cheating provers' replies one at a time shows that the execution with the cheating provers is accepted with probability at least $1 - \epsilon_c - k\eta$, even when the verifier chooses its questions adaptively.
- If y is uniform, then y is not in the image of PRG with high probability. Since $x \notin L$, statistical soundness of the NIWI implies that X has no valid NIWI proof, so statistical soundness of the multiprover proof bounds the acceptance probability of any provers.

Running the verifier with the cheating provers therefore distinguishes $\mathsf{PRG}(s)$ from a uniform string with advantage at least $\Delta - k\eta$, where $\Delta = 1 - \epsilon_c - \epsilon_s$ (Theorem 9.1). Thus, the security of the scheme reduces to the PRG alone. This is only meaningful when $\eta < \Delta/k$, so the weak scheme has only inverse-polynomial hardness, which is why we amplify it. The assumption $\mathcal{A}$ enters only through strong local correctness, that is, through its security against $\mathcal{R}^{\mathcal{P}}$.

Correctness. If $x \in L$, the encryptor's histogram is built from NIWI proofs under s, and the decryptor's from NIWI proofs under w. Even one-bit replies to arbitrary reduction queries may depend on the witness, so we use witness indistinguishability. Consider hybrids that switch the decryptor's proofs from w to s one at a time. Given a single proof, an efficient test can generate all other proofs itself, run both histograms and correlated sampling, and check whether decryption fails; by witness indistinguishability, each switch changes the failure probability only slightly. There are $\mathsf{poly}(T, 2^{\ell_{\mathcal{P}}}, \lambda)$ hybrids, and each test runs in time $\mathsf{poly}(T, 2^{\ell_{\mathcal{P}}}, \lambda)$, where T bounds the running time of the reduction. For one-bit replies and polynomial T, polynomial witness indistinguishability therefore suffices; for SNARGs, we need subexponential witness indistinguishability. In the final hybrid, both histograms are built from independent samples of the same distribution, so by concentration and correlated sampling the two outputs agree with high probability. Note that correctness uses no property of the multiprover proof, and that the tests only run honest encryption and decryption, so witness indistinguishability is needed only against this class (Remark 9.3). In

other words, the computational zero knowledge used for correctness in [LMP25] becomes computational witness indistinguishability, used only for correctness.

Running time. The histograms are over $2^{\ell_{\mathcal{P}}}$ possible replies, so they need only $\mathsf{poly}(T, 2^{\ell_{\mathcal{P}}}, \lambda)$ samples. This gives encryption and decryption time $\mathsf{poly}(T, 2^{\ell_{\mathcal{P}}}, \lambda)$ (Theorem 9.5). For witness encryption and iO, with one-bit replies and a polynomial-time reduction, this is *polynomial*. For SNARGs, with $\ell_{\mathcal{P}} = n^{\epsilon}$, it is $\mathsf{poly}(\lambda) \cdot 2^{O(n^{\epsilon})}$.

Non-uniform reductions. The encryptor runs the verifier, and hence the reduction, on statements X that it samples itself. Thus the construction covers uniform reductions, as well as reductions whose advice depends only on the statement length; in the latter case, the encryptor takes the advice as non-uniform advice, while decryption and key generation remain uniform. It does not cover advice that depends on the statement itself, which the encryptor cannot obtain.

2.4 Implications and Extensions

Both transformations cost time exponential in $\ell_{\mathcal{P}}$, and we apply them with the security parameter chosen as a function $\lambda = \lambda(n)$ of the statement length. Completeness requires $\mathcal{R}^{\mathcal{P}}$, with x and w hardwired, to be efficient in λ, so $n \leq p(\lambda)$ for a fixed polynomial p.

Witness encryption and obfuscation. Here $\ell_{\mathcal{P}} = 1$ for every λ, so we take $\lambda(n) = n$. For a one-query reduction, the coAM verifier then runs in time $\mathsf{poly}(n)$, so $\mathsf{NP} \subseteq \mathsf{coAM}$ and the polynomial hierarchy collapses [BHZ87] (Corollaries 7.1 and 7.3). More generally, a T-time reduction gives verifier time $\mathsf{poly}(T)$, so $T = 2^{o(m)}$ would beat brute force over witnesses of length m (Corollaries 7.2 and 7.4). For a uniform multi-query reduction, Section 2.3 gives polynomial-time witness encryption, and hence public-key encryption as in [GGSW13], whose security reduces to one-way functions alone; the witness-indistinguishability property of the NIWI is used only for correctness (Corollaries 9.7 and 9.8). With advice depending only on the statement length, we instead obtain non-uniform witness encryption and public-key encryption. The construction uses the code of the PRG, so this does not contradict [IR89].

Non-adaptive SNARGs. Here $\ell_{\mathcal{P}}(\lambda, n) = O(\lambda^{\alpha}) \cdot n^{o(1)}$ for a constant $\alpha > 0$, so we want λ small, but only polynomially smaller than n. For a constant $\epsilon > 0$, we take $\lambda(n) = \lceil n^{1/k} \rceil$ for a constant integer $k > \alpha/\epsilon$, so that $n \leq \lambda^k$ and $\ell_{\mathcal{P}} \leq n^{\epsilon}$ for all sufficiently large n. For a one-query reduction, this gives $L \in \mathsf{coAM}[2^{n^{\epsilon}}]$ with polynomial communication, for every constant $\epsilon > 0$ (Corollary 7.5). For a uniform multi-query reduction, Section 2.3 gives witness encryption with encryption and decryption time $\mathsf{poly}(\lambda) \cdot 2^{O(n^{\epsilon})}$ and polynomial-size ciphertexts, assuming subexponentially secure one-way functions and NIWIs. For public-key encryption, we encrypt under the statement that the public key y is a PRG image, with a seed of length $\kappa_0 = (\log \lambda)^{\Theta(1/\gamma)}$ for a γ-subexponentially secure PRG. Taking $\epsilon = \Theta(\gamma)$ sufficiently small gives $|y|^{\epsilon} = o(\log \lambda)$, so encryption and decryption run in time $\mathsf{poly}(\lambda)$, while $\kappa_0^{\gamma} = \omega(\log \lambda)$ keeps the PRG secure against all polynomial-size adversaries (Corollaries 9.6 and 9.9). As before, advice depending only on the statement length gives non-uniform schemes.

Adaptively sound SNARGs. Here the cheating prover chooses x after seeing $\mathsf{pk} \leftarrow \mathcal{V}_1(1^{\lambda}, 1^n)$, so the reduction and its advice depend only on λ and n, and the first message of $\mathcal{V}'$ is independent of x. With the same choice of $\lambda(n)$, fixing the verifier's randomness as non-uniform advice and brute-forcing over the $2^{\ell_{\mathcal{P}}}$ replies gives $L \in \mathsf{SIZE}[2^{n^{\epsilon}}]$ for one-query reductions (Corollary 7.8). For multi-query reductions, we adapt the argument of Gentry and Wichs [GW11], replacing their hard subset-membership problem by hard distributions on true and false statements obtained from worst-case hardness via a minimax argument [LY94]. Adaptivity lets the cheating prover

use a fresh statement on every query, which this argument requires. We give an overview in Appendix A.1.

3 Related Work

We compare our results with known constructions of obfuscation and SNARGs in Appendix B.

Lower bounds for obfuscation. Prior lower bounds for iO rule out constructions that use specific underlying primitives in a restricted way. Mahmoody et al. [MMN+16] rule out fully black-box constructions of perfectly correct iO from primitives realizable in the random oracle model, and Garg, Mahmoody, and Mohammed [GMM17a, GMM17b] extend such separations to monolithic constructions from primitives including witness encryption, fully homomorphic encryption, and single-key functional encryption. In contrast, our lower bounds apply to arbitrary efficiently falsifiable assumptions and place no restriction on the construction; the only restriction is that the reduction uses the adversary as a black box. Finally, the lower bound of Chen and Jin [CJ25] already yields a barrier for succinct iO, where the size of the obfuscated circuit is independent of the input length, and succinct WE, where the ciphertext size is independent of the witness length. Our lower bounds impose no such succinctness requirement.

From arguments to proofs. Rothblum and Vadhan [RV10] show that an interactive argument with a black-box soundness reduction to a falsifiable assumption can be turned into a probabilistically-checkable proof (PCP). The PCP records the prover's response to every possible verifier message, and its verifier runs both the argument verifier and the reduction against this fixed strategy. As in our work, the assumption only needs to be secure against the reduction run with the honest prover. Since PCPs exist unconditionally, their result suggests that PCPs are inherent in such arguments. However, their work does not establish any lower bound.

Black-box reductions and coAM. There is a long history of deriving interactive proofs, and hence complexity-theoretic barriers, from black-box reductions. Examples include barriers to worst-case to average-case reductions for NP [FF93, BT03], and to basing one-way functions and other primitives on NP-hardness [AGGM06, HMX10, BB15]. In these works, the reduction is from an NP-hard problem to breaking a primitive, and the main difficulty is for the AM verifier to emulate a specific inefficient oracle, such as an inverter, using set-size estimation protocols. In our setting, the reduction is from the soundness of an argument to an assumption, and its guarantee holds for *every* oracle, so the cheating prover needs no emulation. Instead, the security of the assumption enters through completeness, and set-size estimation enters only in the complementation step, through [GVW02]. As in these works, a restriction on how the reduction queries its oracle plays a central role: non-adaptivity there, and a single query here.

4 Preliminaries

4.1 Cryptographic Definitions and Primitives

Throughout, an NP relation $\mathcal{R}$ is polynomial-time decidable. For each such relation, fix a polynomial q such that every $(x, w) \in \mathcal{R}$ satisfies $m = |w| \leq q(n)$, where $n = |x|$.

4.1.1 Games and Security

We adapt the game framework of Canetti, Lin, and Pass [CLP13].

Definition 4.1 (Security game, adapted from [CLP13, Definition 5])**.** *A* security game $(\mathcal{C}, c)$ *consists of a probabilistic challenger* $\mathcal{C}$ *and a constant* baseline $c \in [0, 1)$*. On common input* 1^λ*, the challenger interacts with an adversary* $\mathcal{A}$ *and outputs either* win *or* lose*. We say that* $\mathcal{A}$ wins *if the output is* win*. For a function* $T_{\mathcal{C}}$*, the game is* $T_{\mathcal{C}}$-time falsifiable *if, on input* 1^λ*, the challenger runs in time at most* $T_{\mathcal{C}}(\lambda)$*, and* efficiently falsifiable *if* $T_{\mathcal{C}}$ *is polynomial.*

The advantage function *of an adversary family* $\mathcal{A} = \{\mathcal{A}_\lambda\}$ *is*

$$\delta(\lambda) = \mathsf{Adv}_{(\mathcal{C},c)}(\mathcal{A}_\lambda; \lambda) = \max\left\{0,\ \Pr\left[\mathcal{A}_\lambda(1^\lambda) \textit{ wins } \mathcal{C}(1^\lambda)\right] - c\right\}.$$

The probability is over the challenger and adversary coins. The family $\mathcal{A}$ breaks *the game if* δ *is nonnegligible: there is a polynomial* p *such that* $\delta(\lambda) > 1/p(\lambda)$ *for infinitely many* λ*. More generally, for a bound* Γ*, the family* Γ-breaks *the game if there is a polynomial* p *such that* $\delta(\lambda) > 1/p(\Gamma(\lambda))$ *for infinitely many* λ*.*

Definition 4.2 (Γ-security)**.** *Let* $\Gamma : \mathbb{N} \to \mathbb{N}$ *be a positive security bound. A game* $(\mathcal{C}, c)$ *has* Γ-security *if no non-uniform adversary family* $\mathcal{A} = \{\mathcal{A}_\lambda\}$ *with circuit size* $|\mathcal{A}_\lambda| \le \Gamma(\lambda)$ Γ*-breaks it (Definition 4.1).*

Definition 4.3 (Polynomial security)**.** *A game is* polynomially secure *if it has* Γ*-security for every polynomial* Γ*.*

Definition 4.4 (Subexponential security)**.** *A game is* subexponentially secure *if there is a constant* $\epsilon > 0$ *such that it has* Γ*-security for every bound* $\Gamma(\lambda) = 2^{O(\lambda^\epsilon)}$*.*

A *game-based assumption* asserts the Γ-security of a game $(\mathcal{C}, c)$ [CLP13, Definition 6]. When no security bound is specified, the assumption asserts polynomial security.

We call an assumption *falsifiable* if its game is efficiently falsifiable [GW11]. Throughout, we assume that all resource bounds (running times, message lengths, randomness, and round and prover counts) are time-constructible, and that baselines, constants, and error bounds are efficiently computable; we do not repeat these conditions. Unlike [CLP13, Definition 6], we do not require an efficient trivial strategy attaining the baseline c; the arguments below do not use this requirement.

Definition 4.5 (Extended-game-based security, adapted from [CLP13, Definition 15])**.** *Fix an implementation* Π*. Its security is* extended-game-based *if it is specified by a family*

$$\{(\mathcal{C}_x, c)\}_{x \in X}, \qquad X \subseteq \{0,1\}^*,$$

with a common baseline $c \in [0, 1)$*. The challengers are implemented by a single algorithm with* x *as a public parameter fixed before the interaction. Its running time is bounded by a single polynomial in* λ*,* $|x|$*, and the total length of the messages it receives, for every* $x \in \{0,1\}^*$*. For an index sequence* $\{x_\lambda \in X\}$ *and an adversary family* $\mathcal{A} = \{\mathcal{A}_\lambda\}$*, define*

$$\delta(\lambda) = \mathsf{Adv}_{(\mathcal{C}_{x_\lambda},c)}(\mathcal{A}_\lambda; \lambda).$$

The adversary family breaks Π *if* δ *is nonnegligible for some polynomially bounded index sequence. The family has* Γ-security *if, for every polynomial* N*, every sequence* $\{x_\lambda \in X\}$ *with* $|x_\lambda| \le N(\lambda)$*, and every adversary family of size at most* $\Gamma(\lambda)$*, the function* δ *satisfies the bound in Definition 4.2. Polynomial and subexponential security are defined from this condition as in Definitions 4.3 and 4.4.*

4.1.2 Two-Message Laconic Arguments

Syntax. A two-message laconic argument for an NP relation $\mathcal{R} \subseteq \{0,1\}^* \times \{0,1\}^*$, with associated language $L_{\mathcal{R}} = \{x : \exists w \ (x,w) \in \mathcal{R}\}$, is a pair $(\mathcal{V} = (\mathcal{V}_1, \mathcal{V}_2), \mathcal{P})$ of a verifier and a prover, specified by the following polynomial time algorithms:

$\mathcal{V}_1(1^\lambda, x) \to (\mathsf{pk}, \mathsf{vk})$**.** This is a probabilistic algorithm that takes as input a security parameter 1^λ and an instance x. It outputs a first message pk, which is sent to the prover, and a private verifier state vk.

$\mathcal{P}(\mathsf{pk}, x, w) \to \pi$**.** This is a probabilistic algorithm that takes as input a first message pk, an instance x and corresponding witness w. It outputs a prover message π.

$\mathcal{V}_2(\mathsf{vk}, x, \pi) \to \{0,1\}$**.** This is a deterministic algorithm that takes as input a verifier state vk, an instance x and a prover message π. It outputs a bit (1 to accept, 0 to reject).

Definition 4.6 (Two-message laconic argument, adapted from [BIOW20, JKLM25])**.** *A two-message laconic argument $(\mathcal{V} = (\mathcal{V}_1, \mathcal{V}_2), \mathcal{P})$ for an* NP *relation $\mathcal{R} \subseteq \{0,1\}^* \times \{0,1\}^*$, with associated language $L_{\mathcal{R}}$, is required to satisfy the following properties:*

Completeness. *For every polynomial p, there exists a negligible function κ such that, for every $\lambda, n, m \in \mathbb{N}$ with $n \le p(\lambda)$ and $m \le q(n)$, and every $x \in \{0,1\}^n$, $w \in \{0,1\}^m$ with $(x,w) \in \mathcal{R}$,*

$$\Pr\left[\mathcal{V}_2(\mathsf{vk}, x, \pi) = 1 \ : \ \begin{array}{l}(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x) \\ \pi \leftarrow \mathcal{P}(\mathsf{pk}, x, w)\end{array}\right] \ge 1 - \kappa(\lambda).$$

Succinctness. *The length of the prover message is given by $|\pi| = \mathsf{poly}(\lambda)(|x|+|w|)^{o(1)}$. The size of $(\mathsf{pk}, \mathsf{vk})$ may be polynomial in both λ and $|x|$.*

Soundness. *We consider two versions of soundness.*

- ***Search security.*** *For every polynomial p, every sequence of instances $\{x_\lambda \notin L_{\mathcal{R}}\}$ with $|x_\lambda| \le p(\lambda)$, and every non-uniform polynomial-size prover $\mathcal{P}^* = \{\mathcal{P}^*_\lambda\}$, there exists a negligible function ν such that*

$$\Pr\left[\mathcal{V}_2(\mathsf{vk}, x_\lambda, \pi) = 1 \ : \ \begin{array}{l}(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x_\lambda) \\ \pi \leftarrow \mathcal{P}^*_\lambda(\mathsf{pk})\end{array}\right] \le \nu(\lambda).$$

  We can express this as the hardness of finding an accepting prover message for a false statement. Let $X = \{x : x \notin L_{\mathcal{R}}\}$, and let $\mathcal{C}_x$ sample $(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x)$, send pk, receive π, and output win precisely when $\mathcal{V}_2(\mathsf{vk}, x, \pi) = 1$. Search security is equivalent to polynomial security of the family $\{(\mathcal{C}_x, 0)\}_{x \in X}$ (Definition 4.5).

- ***Indistinguishability security.*** *In the designated-verifier setting, an indistinguishability formulation can also be natural. For example, if π is a single bit and the verifier accepts precisely when it equals a hidden uniform bit, soundness asks that no efficient prover predict this bit with probability greater than $1/2$ plus a negligible function. More generally, we allow a constant, efficiently computable soundness baseline $c_0 \in [0,1)$ by requiring polynomial security of $\{(\mathcal{C}_x, c_0)\}_{x \in X}$.*

For a possibly unbounded adversary, its advantage in the game $(\mathcal{C}_x, c_0)$ is its non-adaptive soundness advantage at (λ, x). Unless specified otherwise, $c_0 = 0$.

4.1.3 Succinct Non-Interactive Arguments (SNARGs)

Syntax. A SNARG for an NP relation $\mathcal{R} \subseteq \{0,1\}^* \times \{0,1\}^*$, with associated language $L_{\mathcal{R}} = \{x : \exists w\ (x,w) \in \mathcal{R}\}$, is a triple $\Pi = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ of polynomial-time algorithms:

$\mathcal{G}(1^\lambda, 1^n) \to (\mathsf{pk}, \mathsf{vk})$**.** This probabilistic algorithm outputs a public proving key pk and a private verification key vk.

$\mathcal{P}(\mathsf{pk}, x, w) \to \pi$**.** This probabilistic algorithm takes a proving key, an instance x, and a witness w. It outputs a proof π.

$\mathcal{V}(\mathsf{vk}, x, \pi) \to \{0,1\}$**.** This deterministic algorithm takes a verification key, an instance x, and a proof π. It outputs a bit, with 1 indicating acceptance.

The setup may depend on $n = |x|$, but not on x. This syntax includes designated-verifier SNARGs; public verification is the special case $\mathsf{pk} = \mathsf{vk} = \mathsf{crs}$. Length-independent setup is allowed by ignoring 1^n.

Definition 4.7 (SNARG, adapted from [CJ25, GW11])**.** *A SNARG for an* NP *relation $\mathcal{R} \subseteq \{0,1\}^* \times \{0,1\}^*$, with associated language $L_{\mathcal{R}}$, is required to satisfy the following properties. We use the completeness and succinctness conditions of Gentry and Wichs [GW11, Definition 2.1].*

Completeness. *For every polynomial p, there exists a negligible function κ such that, for every $\lambda, n, m \in \mathbb{N}$ with $n \le p(\lambda)$ and $m \le q(n)$, and every $x \in \{0,1\}^n$, $w \in \{0,1\}^m$ with $(x,w) \in \mathcal{R}$,*

$$\Pr\left[\mathcal{V}(\mathsf{vk}, x, \pi) = 1 \ : \ \begin{array}{c}(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{G}(1^\lambda, 1^n) \\ \pi \leftarrow \mathcal{P}(\mathsf{pk}, x, w)\end{array}\right] \ge 1 - \kappa(\lambda).$$

Succinctness. *The length of a proof is given by $|\pi| = \mathsf{poly}(\lambda)(|x| + |w|)^{o(1)}$, as in Definition 4.6.*

Non-Adaptive Soundness. *For every polynomial p, every sequence of instances $\{x_\lambda \notin L_{\mathcal{R}}\}$ with $|x_\lambda| \le p(\lambda)$, and every non-uniform polynomial-size prover $\mathcal{P}^* = \{\mathcal{P}^*_\lambda\}$, there exists a negligible function ν such that*

$$\Pr\left[\mathcal{V}(\mathsf{vk}, x_\lambda, \pi) = 1 \ : \ \begin{array}{c}(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{G}(1^\lambda, 1^{|x_\lambda|}) \\ \pi \leftarrow \mathcal{P}^*_\lambda(\mathsf{pk})\end{array}\right] \le \nu(\lambda).$$

Adaptive Soundness. *For every polynomial p, every sequence of lengths $\{n_\lambda\}$ with $n_\lambda \le p(\lambda)$, and every non-uniform polynomial-size prover $\mathcal{P}^* = \{\mathcal{P}^*_\lambda\}$, there exists a negligible function ν such that*

$$\Pr\left[\begin{array}{c}|x| = n_\lambda,\ x \notin L_{\mathcal{R}}, \\ \mathcal{V}(\mathsf{vk}, x, \pi) = 1\end{array} \ : \ \begin{array}{c}(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{G}(1^\lambda, 1^{n_\lambda}) \\ (x, \pi) \leftarrow \mathcal{P}^*_\lambda(\mathsf{pk})\end{array}\right] \le \nu(\lambda).$$

Soundness as security games. By Definition 4.5, we can rewrite both soundness notions as the polynomial security of a family of games, as follows.

- *Non-adaptive soundness.* Let $\mathcal{C}_x$ sample $(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{G}(1^\lambda, 1^{|x|})$, send pk, receive π, and output win precisely when $\mathcal{V}(\mathsf{vk}, x, \pi) = 1$. Non-adaptive soundness is equivalent to polynomial security of the family $\{(\mathcal{C}_x, 0)\}_{x \in X}$, with $X = \{x : x \notin L_{\mathcal{R}}\}$. For a possibly unbounded adversary, its advantage in this game is its *non-adaptive soundness advantage* at (λ, x).

- *Adaptive soundness.* Let $\mathcal{C}_{1^n}$ sample $(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{G}(1^\lambda, 1^n)$, send pk, receive (x, π), and output win precisely when $|x| = n$, $x \notin L_{\mathcal{R}}$, and $\mathcal{V}(\mathsf{vk}, x, \pi) = 1$. Adaptive soundness is equivalent to polynomial security of the family of possibly inefficient games $\{(\mathcal{C}_{1^n}, 0)\}_{1^n \in X}$, with $X = \{1^n : n \in \mathbb{N}\}$. For a possibly unbounded adversary, its advantage in this game is its *adaptive soundness advantage* at (λ, n).

In both cases, the index sequence is fixed before setup and may be hardwired into the adversary. In the non-adaptive game, falsity is imposed through X and is not checked by the challenger. In the adaptive game, the instance is chosen by the adversary after setup, so the challenger must check $x \notin L_{\mathcal{R}}$; this check need not be efficient.

4.1.4 Witness Encryption

Syntax. Let $\mathcal{R}$ be an NP relation with language L, and let $\ell = \ell(\lambda)$ be a polynomial message length. A *witness encryption scheme* for $\mathcal{R}$ is a pair $(\mathsf{Enc}, \mathsf{Dec})$ of probabilistic polynomial-time algorithms with the following syntax:

$\mathsf{Enc}(1^\lambda, x, \mu) \to \mathsf{ct}$**.** This algorithm encrypts a message $\mu \in \{0,1\}^{\ell(\lambda)}$ under a statement x and outputs a ciphertext ct.

$\mathsf{Dec}(\mathsf{ct}, w) \to \mu'$**.** This algorithm outputs a message $\mu' \in \{0,1\}^{\ell(\lambda)}$ or a failure symbol $\perp$.

Definition 4.8 (Witness encryption [GGSW13])**.** *A witness encryption scheme* $(\mathsf{Enc}, \mathsf{Dec})$ *for* $\mathcal{R}$*, with message length* $\ell(\lambda)$*, is required to satisfy the following properties:*

Correctness. *There is a negligible function* κ *such that, for all* λ*, all* $(x, w) \in \mathcal{R}$*, and all* $\mu \in \{0,1\}^{\ell(\lambda)}$*,*[11]

$$\Pr\big[\mathsf{Dec}(\mathsf{Enc}(1^\lambda, x, \mu), w) = \mu\big] \geq 1 - \kappa(\lambda).$$

Security. *On input* 1^λ *and with public parameter* x*, the challenger* $\mathcal{C}_x^{\mathsf{WE}}$ *receives two messages* $\mu_0, \mu_1 \in \{0,1\}^{\ell(\lambda)}$ *from the adversary, samples* $b \leftarrow \{0,1\}$*, sends* $\mathsf{Enc}(1^\lambda, x, \mu_b)$*, receives a bit* b'*, and outputs* win *precisely when* $b' = b$*. For every polynomial* p *and every non-uniform polynomial-size adversary* $\mathcal{A}$*, there exists a negligible function* ν *such that, for all* λ *and every* $x \notin L$ *with* $|x| \leq p(\lambda)$*,*

$$\Pr\big[\mathcal{A}(1^\lambda, x) \text{ wins } \mathcal{C}_x^{\mathsf{WE}}(1^\lambda)\big] \leq \frac{1}{2} + \nu(\lambda).$$

Security as a family of games. Equivalently, the family $\{(\mathcal{C}_x^{\mathsf{WE}}, 1/2)\}_{x \in X}$ with $X = \{x : x \notin L\}$ is polynomially secure (Definition 4.5).

For one-bit messages, we also use the equivalent two-message game with the message pair fixed to $(0, 1)$. The challenger samples $b \leftarrow \{0,1\}$, sends $\mathsf{Enc}(1^\lambda, x, b)$, receives a bit b', and wins if $b' = b$; its baseline is $1/2$. A reduction for the chosen-message game can be restricted to adversaries that always choose $(0, 1)$: the reduction supplies this pair internally and forwards only the challenge ciphertext to the response oracle. Our query count refers to calls to this stateless ciphertext-to-bit oracle.

As for non-adaptive soundness, falsity of x is imposed through the index set X and is not checked by the challenger. We say that WE exists for NP if every NP relation has a witness encryption scheme. Γ-security and subexponential security are defined analogously.

[11] Correctness error can be reduced to $T^{-\omega(1)}$ by the standard parallel repetition and majority vote trick, which preserves one-query reductions.

We also consider witness encryption with encryption and decryption time $T(n, \lambda)$, where $n = |x|$. This satisfies the same correctness and security conditions, with the polynomial-time requirement replaced by the stated bound and the security games allowed to be inefficient.

We also use the following weak form of witness encryption, in which the encryptor samples a random message that only needs to be hard to compute.

Syntax. Let $\mathcal{R}$ be an NP relation with language L. A *weak witness encryption scheme* for $\mathcal{R}$ is a pair $(\mathsf{Enc}', \mathsf{Dec}')$ of probabilistic algorithms with the following syntax:

$\mathsf{Enc}'(1^\lambda, x) \to (\mathsf{ct}, \mu)$**.** On a statement x of length n, this algorithm outputs a ciphertext ct and a string μ.

$\mathsf{Dec}'(\mathsf{ct}, w) \to \mu'$**.** This algorithm outputs a string μ'.

Both algorithms run in time at most $T(n, \lambda)$, which may be superpolynomial.

Definition 4.9 (Weak witness encryption, adapted from [LMP25, Definition 4.1])**.** *Let $\alpha, \beta \geq 1$ be functions of (n, λ). A weak witness encryption scheme $(\mathsf{Enc}', \mathsf{Dec}')$ for $\mathcal{R}$ with correctness $1 - 1/\alpha$, hardness $1/\beta$, and running time $t(n, \lambda)$ is required to satisfy the following properties:*

Correctness. *For all sufficiently large λ and all $(x, w) \in \mathcal{R}$, with $n = |x|$,*

$$\Pr\left[\mathsf{Dec}'(\mathsf{ct}, w) = \mu \ : \ (\mathsf{ct}, \mu) \leftarrow \mathsf{Enc}'(1^\lambda, x)\right] \geq 1 - \frac{1}{\alpha(n, \lambda)}.$$

Hardness. *For every non-uniform adversary family $\{\mathcal{A}_{n,\lambda}\}$ of size $\mathsf{poly}(n, \lambda)$, for all sufficiently large λ, all n, and all $x \notin L$ of length n,*

$$\Pr\left[\mathcal{A}_{n,\lambda}(\mathsf{ct}) \neq \mu \ : \ (\mathsf{ct}, \mu) \leftarrow \mathsf{Enc}'(1^\lambda, x)\right] \geq \frac{1}{\beta(n, \lambda)}.$$

Hardness against a larger circuit-size bound means that the same condition holds for adversaries of that size.

Theorem 4.10 (Adapted from [LMP25, Proposition 4.2])**.** *Let $B(n, \lambda) \geq n + \lambda$. Suppose that there is a weak witness encryption scheme for $\mathcal{R}$ with running time $t(n, \lambda) \geq B(n, \lambda)$, correctness $1 - 1/\alpha(n, \lambda)$, and hardness $1/\beta(n, \lambda)$, where $\alpha, \beta \leq \mathsf{poly}(B)$ and*

$$\alpha(n, \lambda) \geq 4B(n, \lambda)\beta(n, \lambda).$$

Assume that ciphertexts and sampled strings have length $\mathsf{poly}(B)$. If hardness holds against adversaries of size $\mathsf{poly}(t(n, \lambda))$, then there is a witness encryption scheme satisfying the correctness and security conditions of Definition 4.8, with running time $\mathsf{poly}(t(n, \lambda))$, whose ciphertexts consist of $\mathsf{poly}(B)$ ciphertexts of the weak scheme and $\mathsf{poly}(B)$ additional bits. Moreover, a polynomial-size distinguisher breaking its security on $x \notin L$ yields, in a black-box way, an adversary of size $\mathsf{poly}(t(n, \lambda))$ violating hardness of the weak scheme on x.

Proof. Apply the amplification of [LMP25] with parameter B. Parallel repetition $\lfloor B\beta \rfloor$ times has correctness at least $3/4$ by a union bound. Hardness amplification for weakly verifiable puzzles [CHS05] makes recovering the repeated secret harder than any inverse polynomial in B: a predictor with such success gives a weak-scheme adversary in time $\mathsf{poly}(t, B)$. The repeated secret has length $\mathsf{poly}(B)$, so Goldreich–Levin reconstruction [GL89] also loses only a polynomial in B. Finally, $O(B)$ repetitions with majority decoding give correctness error $e^{-\Omega(B)}$; a security hybrid loses only $O(B)$. The final ciphertext consists of $O(B^2\beta) \leq \mathsf{poly}(B)$ weak ciphertexts and the $\mathsf{poly}(B)$-bit Goldreich–Levin randomness and masked bit. All reconstruction algorithms use the adversary as a black box and have size $\mathsf{poly}(t, B) = \mathsf{poly}(t)$. □

Taking $B = n + \lambda$ gives the usual version with polynomially bounded α, β and output lengths. Allowing larger B also covers weak schemes whose hardness $1/\beta$ is smaller than inverse polynomial in (n, λ).

4.1.5 Indistinguishability Obfuscation

Syntax. An *indistinguishability obfuscator* for circuits is a probabilistic polynomial-time algorithm iO that takes as input 1^λ and a circuit C, and outputs a circuit $\tilde{C}$.

Definition 4.11 (Indistinguishability obfuscation [BGI+01, GGH+13])**.** *An indistinguishability obfuscator* iO *for circuits is required to satisfy the following properties:*

Correctness. *There is a negligible function κ such that, for all λ, all circuits C, and all inputs u to C,*[12]

$$\Pr\big[\tilde{C}(u) = C(u) \ : \ \tilde{C} \leftarrow \mathsf{iO}(1^\lambda, C)\big] \geq 1 - \kappa(\lambda).$$

Security. *On input 1^λ and with public parameter (C_0, C_1), the challenger $\mathcal{C}^{\mathsf{iO}}_{(C_0,C_1)}$ samples $b \leftarrow \{0,1\}$, sends $\mathsf{iO}(1^\lambda, C_b)$, receives a bit b', and outputs* win *precisely when $b' = b$. For every polynomial p and every non-uniform polynomial-size adversary $\mathcal{A}$, there exists a negligible function ν such that, for all λ and every pair of functionally equivalent circuits C_0, C_1 with the same input length and $|C_0| = |C_1| \leq p(\lambda)$,*

$$\Pr\big[\mathcal{A}(1^\lambda, C_0, C_1) \text{ wins } \mathcal{C}^{\mathsf{iO}}_{(C_0,C_1)}(1^\lambda)\big] \leq \frac{1}{2} + \nu(\lambda).$$

Security as a family of games. Equivalently, the family $\{(\mathcal{C}^{\mathsf{iO}}_{(C_0,C_1)}, 1/2)\}_{(C_0,C_1)\in X}$ is polynomially secure (Definition 4.5), where X is the set of pairs of functionally equivalent circuits of equal size and equal input length.

As before, functional equivalence is imposed through the index set X and is not checked by the challenger, which need not be efficient to decide. The index (C_0, C_1) is public. Γ-security and subexponential security are defined analogously.

4.1.6 Black-Box Reductions

We use the adversary-independent advice model of Chen and Jin [CJ25, Definitions 3.8 and 3.10], with a public index and an explicit security bound. We restrict the games whose security is being proved to *two-message games*: the challenger sends one challenge and receives one response. Thus a source adversary is a map from challenges to responses. Each oracle call to a randomized adversary uses fresh independent coins, with no state shared between calls. The reduction sees the responses but has no access to the adversary's code or random tape through its oracle. The game $(\mathcal{C}, c)$ to which security is reduced may have several rounds.

Let $\{(\mathcal{D}_x, c_0)\}_{x\in X}$ be a family of two-message games as in Definition 4.5. Source challengers may be inefficient. The target $(\mathcal{C}, c)$ is a single game.

Syntax. Let $T(\lambda, n) \geq \lambda + n$. A non-uniform T-time black-box reduction $(\mathcal{R}, F)$ from the security of $\{(\mathcal{D}_x, c_0)\}_{x\in X}$ to the security of $(\mathcal{C}, c)$ consists of a uniform oracle machine $\mathcal{R}$ and a possibly randomized, computationally unbounded advice mapping F, with $|F(1^\lambda, x; r)| \leq T(\lambda, |x|)$ for every λ, $x \in X$, and choice of coins r, together with a deterministic *target-parameter map* f such

[12] As for witness encryption, correctness error can be reduced to $T^{-\omega(1)}$ while preserving one-query reductions.

that $f(1^\lambda, n)$ is computable in time $T(\lambda, n)$. The reduction runs the target challenger on $1^{\lambda'}$, where $\lambda' = f(1^\lambda, |x|)$.

The advice is sampled once, independently of all other coins. All oracle calls use the source parameter λ and the fixed source index; only the target challenger runs at λ'. The machine $\mathcal{R}$ is defined on every binary index x and advice string z. The total running time of one target-game execution, including computing $f(1^\lambda, |x|)$ and the computation of the target challenger $\mathcal{C}(1^{\lambda'})$, is at most $T(\lambda, |x|)$ whenever $|z| \le T(\lambda, |x|)$, for every oracle and even when $x \notin X$, counting oracle communication but excluding the oracle's internal computation. In particular, the target game is run as a T-time falsifiable game, and the whole execution can be simulated in time $\mathsf{poly}(T(\lambda, |x|))$ given oracle answers. The pair $(\mathcal{R}, F)$ is fixed before the index and adversary families; in particular, F does not receive a description of the adversary.

We use the following form of hardness for the target game, which accounts for the target-parameter map.

Definition 4.12 (Target hardness)**.** *Let $T(\lambda, n) \ge \lambda + n$, let f be a target-parameter map, and let p be a polynomial bounding the allowed index lengths. A game $(\mathcal{C}, c)$ is* T-hard at f *if, for every constant $K > 0$, there is a security bound Γ' such that $(\mathcal{C}, c)$ has Γ'-security and*

$$\Gamma'\big(f(1^\lambda, n)\big) \ge T(\lambda, n)^K$$

for all sufficiently large λ and all $n \le p(\lambda)$.

For $f(1^\lambda, n) = \lambda$ and $T(\lambda, n) = \mathsf{poly}(\lambda, n)$, this is polynomial security; for $T = 2^{O(\lambda^\epsilon)}$, it is subexponential security.

Definition 4.13 (Non-uniform T-time black-box reduction)**.** *Let $(\mathcal{R}, F)$ have the syntax above, and let $\Gamma(\lambda, n) \ge \lambda$ be a bound. For every polynomially bounded sequence $\{x_\lambda \in X\}$ and every possibly unbounded non-uniform adversary family $\mathcal{A} = \{\mathcal{A}_\lambda\}$, define*

$$\delta(\lambda) = \mathsf{Adv}_{(\mathcal{D}_{x_\lambda}, c_0)}(\mathcal{A}_\lambda; \lambda), \qquad \mathcal{A}'_{\lambda,z}(1^\lambda) = \mathcal{R}^{\mathcal{A}_\lambda}(1^\lambda, x_\lambda, z).$$

Write

$$\delta'(\lambda) = \max\Big\{0,\ \Pr\Big[\mathcal{A}'_{\lambda,z_\lambda}(1^\lambda) \text{ wins } \mathcal{C}(1^{f(1^\lambda, |x_\lambda|)}) \ :\ z_\lambda \leftarrow F(1^\lambda, x_\lambda)\Big] - c\Big\}.$$

We say that $(\mathcal{R}, F)$ is a non-uniform T-time black-box reduction to the Γ-hardness of $(\mathcal{C}, c)$ at f *if, whenever δ is nonnegligible, the reduction Γ-breaks $(\mathcal{C}, c)$, run at f, in the sense of Definition 4.1: there is a polynomial p such that $\delta'(\lambda) > 1/p(\Gamma(\lambda, |x_\lambda|))$ for infinitely many λ. When $\Gamma(\lambda, n) = \lambda$, that is, when δ' is nonnegligible whenever δ is, we call $(\mathcal{R}, F)$ a* non-uniform T-time black-box reduction to $(\mathcal{C}, c)$*. The security guarantee is required only for indices in X.*

The definition does not presuppose security of either game. Taking $f(1^\lambda, n) = \lambda$ recovers the same security parameter in both games. The security of $(\mathcal{C}, c)$ needed to cover the running time T is a separate hypothesis.

Our arguments often need security only against one specific family of algorithms, such as the reduction run with the honest prover. We capture this as follows.

Definition 4.14 (Γ-security against a class)**.** *Let f be a target-parameter map, p a polynomial, and $\Gamma(\lambda, n) \ge \lambda$ a bound. Let $\mathcal{D}$ be a probabilistic algorithm that takes 1^λ and an auxiliary input* aux*, together with finite sets $\mathsf{Aux}_{\lambda,n}$ of allowed auxiliary inputs. We say that $(\mathcal{C}, c)$ is* Γ-secure against $\mathcal{D}$ at f *if, for every sequence $\{(n_\lambda, \mathsf{aux}_\lambda)\}$ with $n_\lambda \le p(\lambda)$ and $\mathsf{aux}_\lambda \in \mathsf{Aux}_{\lambda,n_\lambda}$, the family $\{\mathcal{D}(1^\lambda, \mathsf{aux}_\lambda)\}$ does not Γ'-break $(\mathcal{C}, c)$ (Definition 4.1), where $\Gamma'(\lambda) = \Gamma(\lambda, n_\lambda)$ and the challenger is run on $1^{f(1^\lambda, n_\lambda)}$. We mainly use $\Gamma = T$, and then say that $(\mathcal{C}, c)$ is* T-secure *against $\mathcal{D}$ at f.*

Definition 4.15 (Advice description complexity)**.** *Let F be an advice mapping as in Definition 4.13. We say that $F(1^\lambda, x)$ has* description complexity at most s_F *if the map $r \mapsto F(1^\lambda, x; r)$ has a Boolean-circuit description of length at most s_F. We say that F has description complexity $s_F(\lambda, n)$ if this bound holds for every λ and $x \in X$ with $n = |x|$.*

Choose one such circuit description for each (λ, x), independently of the adversary. The reduction can use this description as deterministic advice and generate the sampler's coins privately. This preserves its interaction distribution, query count, and query adaptivity, with polynomial overhead in $T + s_F$. Unless s_F is tracked separately, we use this deterministic-advice representation: $F(1^\lambda, x)$ returns the sampler description, and T bounds both its length and its evaluation cost.

Definition 4.16 (Uniform black-box reduction)**.** *A reduction in Definition 4.13 is* uniform *if F always outputs the empty string.*

Definition 4.17 (Black-box reduction with deterministic advice)**.** *A reduction in Definition 4.13 has* deterministic advice *if F is deterministic.*

Definition 4.18 (k-query black-box reduction)**.** *A reduction is k-*query *if it makes at most $k(\lambda, |x|)$ oracle calls during one target-game execution.*

Definition 4.19 (One-query black-box reduction)**.** *A reduction is* one-query *if it is k-query with $k = 1$.*

Definition 4.20 (Black-box reduction with adaptive queries)**.** *A reduction* with adaptive queries *may choose later oracle queries using earlier oracle replies.*

Definition 4.21 (Black-box reduction with non-adaptive queries)**.** *A reduction has* non-adaptive queries *if it fixes its entire list of oracle queries before receiving any oracle reply.*

For non-adaptive soundness, the index is the statement x, and the advice is $F(1^\lambda, x)$. For adaptive soundness of a SNARG, the index is 1^n, so the advice is $F(1^\lambda, 1^n)$; it cannot depend on the instance subsequently returned by the adversary. Each oracle query is a proving key pk, and the reply is a proof π in the non-adaptive case and a pair (x, π) in the adaptive case. Query adaptivity describes the reduction's oracle calls and is separate from adaptive soundness.

4.1.7 Other Cryptographic Tools

Pseudorandom generators (PRGs). Syntax. Let $n = n(\lambda)$ and $m = m(\lambda)$ be polynomially bounded, efficiently computable length functions with $m(\lambda) > n(\lambda)$. A *PRG* is a family of deterministic functions $\mathsf{PRG}_\lambda : \{0,1\}^{n(\lambda)} \to \{0,1\}^{m(\lambda)}$, where $\mathsf{PRG}_\lambda(s)$ is computable in polynomial time given $(1^\lambda, s)$. We omit the subscript λ when it is clear from context.

Definition 4.22 (PRG [BM82, Yao82])**.** *A PRG* PRG *is* polynomially secure *if, for every non-uniform polynomial-size adversary $\mathcal{A}$, there exists a negligible function ν such that, for all λ, its distinguishing advantage satisfies*

$$\left| \Pr_{s \leftarrow \{0,1\}^{n(\lambda)}} \left[\mathcal{A}(\mathsf{PRG}_\lambda(s)) = 1\right] - \Pr_{y \leftarrow \{0,1\}^{m(\lambda)}} \left[\mathcal{A}(y) = 1\right] \right| \le \nu(\lambda).$$

A uniformly random $y \in \{0,1\}^{m(\lambda)}$ lies in the image of PRG_λ with probability at most $2^{n(\lambda)-m(\lambda)}$. In our applications, we use length-doubling PRGs, with $n(\lambda) = \lambda$ and $m(\lambda) = 2\lambda$, so this probability is at most $2^{-\lambda}$. Subexponentially secure one-way functions give subexponentially secure PRGs [HILL99].

Non-interactive witness-indistinguishable proofs (NIWIs). We consider NIWIs with a (possibly non-uniform) CRS, as they are sufficient for our applications. This makes any NIZK with a statistically sound mode sufficient for our purposes.

Syntax. A *NIWI* for an NP relation $\mathcal{R} \subseteq \{0,1\}^* \times \{0,1\}^*$, with associated language $L_\mathcal{R} = \{x : \exists w\ (x,w) \in \mathcal{R}\}$, is a triple $\Pi = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ of polynomial-time algorithms with the following syntax:

$\mathcal{G}(1^\lambda, 1^n, 1^m) \to \mathsf{crs}$**.** This probabilistic algorithm outputs a common reference string crs.

$\mathcal{P}(\mathsf{crs}, x, w) \to \pi$**.** This probabilistic algorithm takes an instance x and a witness w, and outputs a proof π.

$\mathcal{V}(\mathsf{crs}, x, \pi) \to \{0,1\}$**.** This deterministic algorithm takes an instance x and a proof π, and outputs a bit, with 1 indicating acceptance.

Definition 4.23. *A NIWI $\Pi = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for $\mathcal{R}$ is required to satisfy the following properties:*

Completeness. *For all λ, n, m, and $(x,w) \in (\{0,1\}^n \times \{0,1\}^m) \cap \mathcal{R}$,*

$$\Pr\left[\mathcal{V}(\mathsf{crs}, x, \pi) = 1 \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathcal{G}(1^\lambda, 1^n, 1^m) \\ \pi \leftarrow \mathcal{P}(\mathsf{crs}, x, w) \end{array}\right] = 1.$$

Statistical Soundness. *For every polynomial p, there exists a negligible function κ such that, for all λ, all $n, m \le p(\lambda)$, every $x \in \{0,1\}^n \setminus L_\mathcal{R}$, and every possibly unbounded prover $\mathcal{P}^*$,*

$$\Pr\left[\mathcal{V}(\mathsf{crs}, x, \pi) = 1 \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathcal{G}(1^\lambda, 1^n, 1^m) \\ \pi \leftarrow \mathcal{P}^*(\mathsf{crs}, x) \end{array}\right] \le \kappa(\lambda).$$

Witness Indistinguishability. *For every polynomial p and every non-uniform polynomial-size adversary $\mathcal{A}$, there exists a negligible function κ such that, for all λ, all $n, m \le p(\lambda)$, every $x \in \{0,1\}^n$, and all $w_0, w_1 \in \{0,1\}^m$ with $(x,w_0), (x,w_1) \in \mathcal{R}$,*

$$\left| \Pr\left[\mathcal{A}(\mathsf{crs}, \pi) = 1 \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathcal{G}(1^\lambda, 1^n, 1^m) \\ \pi \leftarrow \mathcal{P}(\mathsf{crs}, x, w_0) \end{array}\right] - \Pr\left[\mathcal{A}(\mathsf{crs}, \pi) = 1 \ : \ \begin{array}{l} \mathsf{crs} \leftarrow \mathcal{G}(1^\lambda, 1^n, 1^m) \\ \pi \leftarrow \mathcal{P}(\mathsf{crs}, x, w_1) \end{array}\right] \right| \le \kappa(\lambda).$$

Statistically sound NIZKs (and so NIWIs) have been instantiated from a few assumptions, including the subgroup decisional assumption [GOS06], decisional linear assumption [GOS12], k-Lin [LPWW20], and learning with errors [PS19, WWW25]. With suitable complexity leveraging, they can achieve subexponential security from subexponential assumptions.

Theorem 4.24 ([GOS06, GOS12, LPWW20, PS19, WWW25])**.** *Assuming either subexponential security of the subgroup decisional assumption over a pairing group, the decisional linear assumption over a pairing group, k-Lin over a pairing group, or learning with errors, there exists a subexponentially secure NIWI for all* NP *relations.*

We use the following non-uniform form of subexponential witness indistinguishability.

Definition 4.25 (Subexponential witness indistinguishability)**.** *For a constant $\epsilon > 0$, a NIWI has ϵ-*subexponential witness indistinguishability *if, for all sufficiently large λ, every statement x, every two witnesses w_0, w_1 for x of the same length, and every non-uniform circuit of size at most 2^{λ^ϵ}, its witness-indistinguishability advantage is at most $2^{-\lambda^\epsilon}$.*

Note that the distinguisher may depend on the statement and both witnesses.

Similarly, for a constant $\epsilon > 0$, a PRG (Definition 4.22) is *ϵ-subexponentially secure* if, for all sufficiently large λ, every non-uniform circuit of size at most 2^{λ^ϵ} has distinguishing advantage at most $2^{-\lambda^\epsilon}$.

Public-key encryption (PKE). Syntax. A *public-key encryption scheme* for one-bit messages is a triple $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ of probabilistic polynomial-time algorithms with the following syntax:

$\mathsf{KeyGen}(1^\lambda) \to (\mathsf{pk}, \mathsf{sk})$**.** This algorithm outputs a public key pk and a secret key sk.

$\mathsf{Enc}(\mathsf{pk}, b) \to \mathsf{ct}$**.** This algorithm encrypts a bit $b \in \{0,1\}$ and outputs a ciphertext ct.

$\mathsf{Dec}(\mathsf{sk}, \mathsf{ct}) \to b'$**.** This algorithm outputs a bit b'.

Definition 4.26 (PKE [GM84])**.** *A public-key encryption scheme $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec})$ for one-bit messages is required to satisfy the following properties:*

Correctness. *There is a negligible function κ such that, for all λ and $b \in \{0,1\}$,*

$$\Pr\left[\mathsf{Dec}(\mathsf{sk}, \mathsf{ct}) = b \;:\; \begin{array}{l}(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda) \\ \mathsf{ct} \leftarrow \mathsf{Enc}(\mathsf{pk}, b)\end{array}\right] \geq 1 - \kappa(\lambda).$$

Security. *For every non-uniform polynomial-size adversary $\mathcal{A}$, there exists a negligible function ν such that, for all λ,*

$$\Big|\Pr\big[\mathcal{A}(\mathsf{pk}, \mathsf{Enc}(\mathsf{pk}, 0)) = 1 \;:\; (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)\big] - \Pr\big[\mathcal{A}(\mathsf{pk}, \mathsf{Enc}(\mathsf{pk}, 1)) = 1 \;:\; (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KeyGen}(1^\lambda)\big]\Big| \leq \nu(\lambda).$$

Here, pk denotes a public key of the encryption scheme, unrelated to the proving key of a SNARG. Encrypting longer messages bit by bit gives public-key encryption for messages of any polynomial length.

4.2 Complexity Theory

A *promise problem* is a pair $(L_{\text{yes}}, L_{\text{no}})$ of disjoint subsets of $\{0,1\}^*$, representing the yes and no instances, respectively. The definitions below extend to promise problems by requiring completeness for $x \in L_{\text{yes}}$ and soundness for $x \in L_{\text{no}}$, with no correctness requirement for $x \notin L_{\text{yes}} \cup L_{\text{no}}$. The complement of $(L_{\text{yes}}, L_{\text{no}})$ is $(L_{\text{no}}, L_{\text{yes}})$.

Minimax theorem. We use von Neumann's minimax theorem for finite two-player zero-sum games. Below, p and q range over distributions on A and B, respectively.

Theorem 4.27 (Minimax theorem [vN28])**.** *Let A, B be finite sets and $M : A \times B \to \mathbb{R}$. Then*

$$\min_p \max_{b \in B} \mathbb{E}_{a \leftarrow p}[M(a,b)] = \max_q \min_{a \in A} \mathbb{E}_{b \leftarrow q}[M(a,b)].$$

4.2.1 Interactive Proofs and Arthur–Merlin Proofs

We recall interactive proofs and Arthur–Merlin proofs. In the definition below, n denotes the common-input length. One round consists of a verifier message followed by a prover reply.

Syntax. An *interactive proof* for a language $L \subseteq \{0,1\}^*$ is a protocol between a computationally unbounded prover $\mathcal{P}$ and a randomized verifier $\mathcal{V}$, both given a common input x. The parties exchange messages, starting with the verifier, after which the verifier outputs a bit. We write $\langle \mathcal{P}, \mathcal{V} \rangle(x)$ for this bit. In a *private-coin* proof, the verifier may keep its randomness secret. A *public-coin* proof, also called an *Arthur–Merlin proof*, restricts every verifier message to a fresh, independent uniform random string, and the verifier's final decision is a deterministic function of the common input and transcript.

Definition 4.28 (Interactive proofs)**.** *Let $\epsilon_c, \epsilon_s : \mathbb{N} \to [0,1]$ satisfy $\epsilon_c(n) + \epsilon_s(n) < 1$. An interactive proof $(\mathcal{P}, \mathcal{V})$ for L with completeness error ϵ_c and soundness error ϵ_s is required to satisfy the following properties:*

Completeness. *For every $x \in L$,*

$$\Pr\big[\langle \mathcal{P}, \mathcal{V} \rangle(x) = 1\big] \geq 1 - \epsilon_c(|x|).$$

Soundness. *For every $x \notin L$ and every computationally unbounded prover $\mathcal{P}^*$,*

$$\Pr\big[\langle \mathcal{P}^*, \mathcal{V} \rangle(x) = 1\big] \leq \epsilon_s(|x|).$$

The probabilities are over the random coins used in the interaction.

We follow the notation of Holenstein and Künzler [HK14, Definition 2.5] and Goldreich, Vadhan, and Wigderson [GVW02, Section 2.1], using named parameters and separate bounds on the total communication from the prover and the verifier. We write

$$\mathsf{IP}\begin{pmatrix} \mathsf{rounds} & = & k(n) \\ \mathsf{verifier\ time} & = & T(n) \\ \mathsf{prover\ comm} & = & \ell_{\mathcal{P}}(n) \\ \mathsf{verifier\ comm} & = & \ell_{\mathcal{V}}(n) \\ \mathsf{verifier\ randomness} & = & r(n) \\ \mathsf{completeness} & \geq & 1 - \epsilon_c(n) \\ \mathsf{soundness} & \leq & \epsilon_s(n) \end{pmatrix}$$

for the class of languages admitting interactive proofs with at most $k(n)$ rounds, verifier running time $O(T(n))$, at most $\ell_{\mathcal{P}}(n)$ prover bits and $\ell_{\mathcal{V}}(n)$ verifier bits in total, at most $r(n)$ verifier random bits, completeness error at most $\epsilon_c(n)$, and soundness error at most $\epsilon_s(n)$. The bounds on verifier time, communication, and randomness hold for every prover strategy and every choice of the verifier's coins. Replacing IP by AM imposes the additional public-coin restriction.

We allow T to be superpolynomial. If the time parameter is omitted, it defaults to $\mathsf{poly}(n)$; omitted communication and randomness bounds default to $O(T(n))$. By default, both completeness and soundness errors are at most $1/3$. Thus, the usual polynomial-time classes are

$$\mathsf{IP} = \mathsf{IP}\big(\mathsf{rounds} = \mathsf{poly}(n)\big), \qquad \mathsf{AM} = \mathsf{AM}\big(\mathsf{rounds} = 1\big).$$

We write $\mathsf{coAM} = \{L : \overline{L} \in \mathsf{AM}\}$; when parameters are specified, coAM denotes the complements of languages in the AM class with the same parameter values.

Parallel repetition and error reduction. Parallel repetition amplifies the completeness–soundness gap of an interactive proof while preserving the number of rounds. We use the threshold version of parallel repetition given by Goldreich, Rothblum, and Skverer [GRS23, full version, Appendix A.2, Theorem A.5], together with the Chernoff bound for completeness.

Lemma 4.29 (Parallel repetition)**.** *Let $\epsilon_c, \epsilon_s : \mathbb{N} \to [0,1]$ satisfy $\epsilon_c(n) + \epsilon_s(n) < 1$, let $\Delta = 1 - \epsilon_c - \epsilon_s$, and let $0 < \eta = \eta(n) < 1/2$. For resource bounds $T, \ell_{\mathcal{P}}, \ell_{\mathcal{V}}, r, k$ and q satisfying*

$$q \geq \frac{2\ln(1/\eta)}{\Delta^2},$$

we have

$$\mathsf{IP}\begin{pmatrix}\text{rounds} & = & k\\ \text{verifier time} & = & T\\ \text{prover comm} & = & \ell_{\mathcal{P}}\\ \text{verifier comm} & = & \ell_{\mathcal{V}}\\ \text{verifier randomness} & = & r\\ \text{completeness} & \geq & 1-\epsilon_c\\ \text{soundness} & \leq & \epsilon_s\end{pmatrix} \subseteq \mathsf{IP}\begin{pmatrix}\text{rounds} & = & k\\ \text{verifier time} & = & qT\\ \text{prover comm} & = & q\ell_{\mathcal{P}}\\ \text{verifier comm} & = & q\ell_{\mathcal{V}}\\ \text{verifier randomness} & = & qr\\ \text{completeness} & \geq & 1-\eta\\ \text{soundness} & \leq & \eta\end{pmatrix}.$$

The same inclusion holds with AM *in place of* IP*. Both inclusions also hold for promise problems.*

From private coins to public coins. Goldwasser and Sipser [GS86] showed that private-coin interactive proofs can be transformed into public-coin proofs with a constant additive increase in the number of rounds. We use the quantitative refinement of Holenstein and Künzler [HK14], which tracks verifier time, message lengths, and randomness separately and allows arbitrary time-constructible verifier time T. Its completeness and soundness requirements are independent of the verifier's randomness complexity. The transformation uses more rounds, but preserves a constant round count for the protocols considered here.

Theorem 4.30 (Private-to-public-coin transformation [HK14, Corollary 4.2(i), full version, with amplification])**.** *Let $k = k(n) \geq 1$ be polynomially bounded, and let $T \geq n$, $\ell_{\mathcal{P}}, \ell_{\mathcal{V}}, r \geq 1$, and k be bounds. Then*

$$\mathsf{IP}\begin{pmatrix}\text{rounds} & = & k\\ \text{verifier time} & = & T\\ \text{prover comm} & = & \ell_{\mathcal{P}}\\ \text{verifier comm} & = & \ell_{\mathcal{V}}\\ \text{verifier randomness} & = & r\end{pmatrix} \subseteq \mathsf{AM}\begin{pmatrix}\text{rounds} & = & 4k+3\\ \text{verifier time} & = & O(kT) + \mathsf{poly}((\ell_{\mathcal{P}}+\ell_{\mathcal{V}}+r)k)\\ \text{prover comm} & = & \mathsf{poly}((\ell_{\mathcal{P}}+\ell_{\mathcal{V}}+r)k)\\ \text{verifier comm} & = & \mathsf{poly}((\ell_{\mathcal{P}}+\ell_{\mathcal{V}}+r)k)\\ \text{verifier randomness} & = & \mathsf{poly}((\ell_{\mathcal{P}}+\ell_{\mathcal{V}}+r)k)\end{pmatrix}.$$

The same inclusion holds for promise problems.

To obtain this formulation, apply Lemma 4.29 with $\eta = 2^{-(k+5)}$ and $q = O(k)$. The amplified proof has completeness at least $3/4$, soundness at most $2^{-(k+5)}$, verifier time $O(kT)$, prover communication $O(k\ell_{\mathcal{P}})$, verifier communication $O(k\ell_{\mathcal{V}})$, and verifier randomness $O(kr)$. Applying [HK14, Corollary 4.2(i)] with $\gamma = 1/12$ gives the theorem. The public-coin verifier invokes the amplified verifier once, and hence runs $O(k)$ copies of the original verifier. In particular, for constant k, the communication and randomness bounds are polynomial in $\ell_{\mathcal{P}} + \ell_{\mathcal{V}} + r$, independently of T.

Reducing the number of rounds. The round-collapse theorem of Babai and Moran [BM88] collapses a constant number of public-coin rounds to one. We use the communication bound of [GVW02, Theorem 2.3], with the verifier's running time tracked through their construction.

Theorem 4.31 (Round-collapse theorem [GVW02, Theorem 2.3 and Appendix B])**.** *For bounds $T = T(n) \geq n$, $\ell_{\mathcal{P}} = \ell_{\mathcal{P}}(n) \geq 1$, $\ell_{\mathcal{V}} = \ell_{\mathcal{V}}(n) \geq 1$, and $k = k(n) \geq 1$,*

$$\mathsf{AM}\begin{pmatrix}\mathsf{rounds} & = & k\\ \mathsf{verifier\ time} & = & T\\ \mathsf{prover\ comm} & = & \ell_{\mathcal{P}}\\ \mathsf{verifier\ comm} & = & \ell_{\mathcal{V}}\\ \mathsf{verifier\ randomness} & = & \ell_{\mathcal{V}}\end{pmatrix} \subseteq \mathsf{AM}\begin{pmatrix}\mathsf{rounds} & = & 1\\ \mathsf{verifier\ time} & = & T\cdot(\ell_{\mathcal{P}}k)^{O(k)}\\ \mathsf{prover\ comm} & = & (\ell_{\mathcal{P}}k)^{O(k)}\\ \mathsf{verifier\ comm} & = & \ell_{\mathcal{V}}\cdot(\ell_{\mathcal{P}}k)^{O(k)}\\ \mathsf{verifier\ randomness} & = & \ell_{\mathcal{V}}\cdot(\ell_{\mathcal{P}}k)^{O(k)}\end{pmatrix}.$$

Both sides have completeness at least $2/3$ and soundness at most $1/3$. In particular, for constant k, the resulting proof has verifier time $T \cdot \mathsf{poly}(\ell_{\mathcal{P}})$, prover-message length $\mathsf{poly}(\ell_{\mathcal{P}})$, and verifier communication and randomness bounded by $\ell_{\mathcal{V}} \cdot \mathsf{poly}(\ell_{\mathcal{P}})$. The same inclusion holds for promise problems.

4.2.2 Multiprover Interactive Proofs

Syntax. Let $L \subseteq \{0,1\}^*$ and let $k = k(n) \geq 1$. A classical k-prover interactive proof for L consists of a uniform randomized verifier $\mathcal{V}$ and computationally unbounded provers $\mathcal{P}_1, \ldots, \mathcal{P}_k$, with common input x of length n. The verifier exchanges messages with each prover over a private channel. A *transcript* records these messages in order, with their senders and recipients. Prover i's *local transcript* consists only of its messages with $\mathcal{V}$, including the current query. The verifier may adapt its queries to the entire transcript; each prover's reply depends only on x, its local transcript, and its coins. The provers may agree on non-uniform strategies and share randomness before the interaction, but cannot communicate during it. They may also use independent private coins; all prover coins are independent of the verifier's coins. We write $\langle \mathcal{P}_1, \ldots, \mathcal{P}_k, \mathcal{V} \rangle(x)$ for the verifier's output bit.

Definition 4.32 (Multiprover interactive proofs, adapted from [Tha23, Definition 8.1])**.** *Let $\epsilon_c, \epsilon_s : \mathbb{N} \to [0,1]$ satisfy $\epsilon_c(n) + \epsilon_s(n) < 1$. A classical k-prover interactive proof for L with completeness error ϵ_c and soundness error ϵ_s is required to satisfy the following properties:*

Completeness. *For every $x \in L$, there exist prover strategies $\mathcal{P}_1, \ldots, \mathcal{P}_k$ such that*

$$\Pr\big[\langle \mathcal{P}_1, \ldots, \mathcal{P}_k, \mathcal{V} \rangle(x) = 1\big] \geq 1 - \epsilon_c(n).$$

Soundness. *For every $x \notin L$ and every tuple of unbounded noncommunicating provers $\mathcal{P}_1^*, \ldots, \mathcal{P}_k^*$,*

$$\Pr\big[\langle \mathcal{P}_1^*, \ldots, \mathcal{P}_k^*, \mathcal{V} \rangle(x) = 1\big] \leq \epsilon_s(n).$$

Remark. In this work, we only consider protocols with at most one round trip between the verifier and each prover: one verifier message followed by one prover reply. The verifier may choose the next prover and its question adaptively from earlier replies; each prover sees only its own question. We call the protocol *non-adaptive* if all questions and their recipients are fixed before any replies are received.

We write

$$\mathsf{MIP}\begin{pmatrix}\text{provers} & = & k(n)\\ \text{verifier time} & = & T(n)\\ \text{ind prover comm} & = & \ell_{\mathcal{P}}(n)\\ \text{verifier comm} & = & \ell_{\mathcal{V}}(n)\\ \text{verifier randomness} & = & r(n)\\ \text{completeness} & \geq & 1-\epsilon_c(n)\\ \text{soundness} & \leq & \epsilon_s(n)\end{pmatrix}$$

for the class of languages admitting such protocols with at most $k(n)$ provers, verifier time $O(T(n))$, at most $\ell_{\mathcal{P}}(n)$ bits sent by each prover, total verifier communication $\ell_{\mathcal{V}}(n)$, and at most $r(n)$ verifier random bits. Resource bounds hold for every prover strategy and every choice of coins. As for IP, T may be superpolynomial; omitted time defaults to $\mathsf{poly}(n)$, omitted communication and randomness bounds to $O(T(n))$, and omitted errors to $1/3$.

4.3 Statistics, Hashing, and Sampling

Notation. For an event A and a sampling experiment E, we write $\Pr\left[\ A\ :\ E\ \right]$ for the probability that A occurs when we run E. For example,

$$\Pr\left[\ A(X)\ :\ X \leftarrow \mathcal{D}\ \right] = \Pr_{X\leftarrow\mathcal{D}}[A(X)].$$

For an event A, its indicator $\mathbf{1}[A]$ equals 1 if A occurs and 0 otherwise. For distributions $\mathcal{D}, \mathcal{D}'$ on a finite set Σ, we write $\|\mathcal{D}-\mathcal{D}'\|_1 = \sum_{\sigma\in\Sigma}|\mathcal{D}(\sigma)-\mathcal{D}'(\sigma)|$, so that their statistical distance is $\mathrm{SD}(\mathcal{D},\mathcal{D}') = \frac{1}{2}\|\mathcal{D}-\mathcal{D}'\|_1$.

Concentration. We recall standard concentration bounds for independent and limited-independence random variables.

Lemma 4.33 (Hoeffding's inequality [Hoe63])**.** *Let $Y_1,\ldots,Y_N$ be independent random variables in $[0,1]$, and put $\overline{\mu} = N^{-1}\sum_i \mathbb{E}[Y_i]$. For every $\eta > 0$,*

$$\Pr\left[\frac{1}{N}\sum_i Y_i \geq \overline{\mu}+\eta\right] \leq e^{-2N\eta^2}, \qquad \Pr\left[\frac{1}{N}\sum_i Y_i \leq \overline{\mu}-\eta\right] \leq e^{-2N\eta^2}.$$

Random variables are *k-wise independent* if every subcollection of at most k variables is mutually independent.

Lemma 4.34 ([BR94, Lemma 2.3, full version])**.** *Let $k \geq 4$ be an even integer, and let $X_1,\ldots,X_v$ be k-wise independent random variables taking values in $[0,1]$. Put $Z = \sum_{j=1}^{v} X_j$ and $\mu = \mathbb{E}[Z]$. For every $\tau > 0$,*

$$\Pr[|Z-\mu| \geq \tau] \leq 8\left(\frac{k\mu+k^2}{\tau^2}\right)^{k/2}.$$

Limited-independence hashing. Syntax. Let $R, b \geq 0$ be integers. An indexed hash family $\mathcal{H} = (h_s)_{s\in\{0,1\}^u}$ consists of functions $h_s : \{0,1\}^R \to \{0,1\}^b$. We call s the *seed* and u the *seed length*. Sampling $h \leftarrow \mathcal{H}$ means sampling its seed uniformly.

Definition 4.35 (k-wise independent hashing [Vad12, Definition 3.31])**.** *Let $k \geq 1$ be an integer. An indexed hash family $\mathcal{H}$ is* k-wise independent *if, for every $1 \leq q \leq \min\{k, 2^R\}$, distinct $x_1, \ldots, x_q \in \{0,1\}^R$, and $y_1, \ldots, y_q \in \{0,1\}^b$,*

$$\Pr_{s \leftarrow \{0,1\}^u}\left[h_s(x_j) = y_j \text{ for all } j \in [q]\right] = 2^{-bq}.$$

In particular, the definition requires the hash values to be uniform as well as independent.

Lemma 4.36 ([Vad12, Corollary 3.34])**.** *For every $R, b, k \geq 1$, there is a k-wise independent family of hashes $h : \{0,1\}^R \to \{0,1\}^b$ with seed length $k \max\{R, b\}$ and evaluation time $\mathsf{poly}(R, b, k)$.*

Sampling. We also use the following elementary fact: a distribution with finite support can be sampled, up to arbitrarily small statistical distance, by a small circuit from uniformly random bits.

Lemma 4.37. *Let $\mathcal{D}$ be a distribution supported on $m \geq 1$ strings $z_1, \ldots, z_m$, each of length at most ℓ, and let $\varepsilon \in (0,1)$. Set $b = \lceil \log_2(m/\varepsilon) \rceil$. Then there is a distribution $\widetilde{\mathcal{D}}$ on $\{z_1, \ldots, z_m\}$, at statistical distance at most ε from $\mathcal{D}$, that is the output distribution of a Boolean circuit with $O(m(\ell + b))$ gates on b uniformly random input bits. In particular, this circuit has a description of length $O(m(\ell + b)\log(m(\ell + b)))$.*

Proof. Let $\alpha_i = \Pr_{\mathcal{D}}[z_i]$. Round each of $\alpha_1, \ldots, \alpha_{m-1}$ down to a multiple of 2^{-b}, obtaining $\widetilde{\alpha}_1, \ldots, \widetilde{\alpha}_{m-1}$, and set $\widetilde{\alpha}_m = 1 - \sum_{i<m} \widetilde{\alpha}_i$. Let $\widetilde{\mathcal{D}}$ assign weight $\widetilde{\alpha}_i$ to z_i. Each of the first $m-1$ weights decreases by less than 2^{-b}, and only z_m gains mass, so the statistical distance is at most $(m-1)2^{-b} \leq \varepsilon$.

All weights of $\widetilde{\mathcal{D}}$ are multiples of 2^{-b}. The circuit reads its b input bits as an integer $u \in \{0, \ldots, 2^b - 1\}$ and outputs z_i for the least i with $u < 2^b \sum_{i' \leq i} \widetilde{\alpha}_{i'}$. Hardwiring the m thresholds and strings, each comparison takes $O(b)$ gates and selecting the output takes $O(m\ell)$ gates. □

We also use empirical distributions, and the correlated sampling procedure of [LMP25], which lets two parties holding close empirical distributions output the same sample using shared randomness. A *histogram* with M samples on a finite set Σ is a function $h : \Sigma \to \{0, \ldots, M\}$ with $\sum_\sigma h(\sigma) = M$. The histogram of M independent samples from $\mathcal{D}$ gives the *empirical distribution* h/M. By Hoeffding's inequality, it is close to $\mathcal{D}$.

Lemma 4.38. *Let $\mathcal{D}$ be a distribution on a finite set Σ, let h be the histogram of M independent samples from $\mathcal{D}$, and let $\eta > 0$. Then, except with probability $2|\Sigma|e^{-2M\eta^2}$, we have that $|h(\sigma)/M - \mathcal{D}(\sigma)| \leq \eta$ for every $\sigma \in \Sigma$. In particular, $\|h/M - \mathcal{D}\|_1 \leq |\Sigma|\eta$.*

Proof. Apply Lemma 4.33 to the indicators of σ for each $\sigma \in \Sigma$, and take a union bound. □

Fix an ordering $\sigma_1, \sigma_2, \ldots$ of Σ. For a histogram h with M samples and $v \in \{0, \ldots, M-1\}$, let $\mathsf{CorSamp}(h, v)$ be the unique σ_t with

$$\sum_{t' < t} h(\sigma_{t'}) \leq v < \sum_{t' \leq t} h(\sigma_{t'}).$$

For uniform v, the output $\mathsf{CorSamp}(h, v)$ is distributed according to h/M. Moreover, two close histograms give the same output for most v.

Lemma 4.39 (Correlated sampling, following [LMP25])**.** *For any two histograms h, h' with M samples on Σ,*

$$\Pr_{v \leftarrow \{0,\ldots,M-1\}}\left[\mathsf{CorSamp}(h, v) \neq \mathsf{CorSamp}(h', v)\right] \leq |\Sigma| \cdot \left\| \frac{h}{M} - \frac{h'}{M} \right\|_1.$$

Proof. Let $F_t = \sum_{t' \le t} h(\sigma_{t'})/M$ and $F'_t = \sum_{t' \le t} h'(\sigma_{t'})/M$. If the two outputs differ, then v/M lies between F_t and F'_t for some t. Since MF_t and MF'_t are integers, this occurs for at most $M|F_t - F'_t|$ values of v. By a union bound, the two outputs differ with probability at most $\sum_t |F_t - F'_t| \le |\Sigma| \cdot \|h/M - h'/M\|_1$. □

5 From Laconic Interactive Proofs to coAM with Polynomial Communication

Goldreich, Vadhan, and Wigderson show that a language with a laconic interactive proof has a public-coin proof for its complement. Here we adapt their result to one-round proofs with an explicit verifier-time parameter and arbitrary completeness and soundness errors.

For one-round interactive proofs, the refinement gives polynomial communication while allowing exponential verifier time.

Theorem 5.1. *Let $T(n) \ge n$ and $\ell_P(n) \ge 1$ be bounds. Let $\epsilon_c(n), \epsilon_s(n) \in [0,1]$ be error bounds with*

$$\Delta(n) = 1 - \epsilon_c(n) - \epsilon_s(n) > 0.$$

If

$$L \in \mathsf{IP} \begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & T \\ \mathsf{prover\ comm} & = & \ell_P \\ \mathsf{completeness} & \ge & 1 - \epsilon_c \\ \mathsf{soundness} & \le & \epsilon_s \end{pmatrix},$$

then

$$L \in \mathsf{coAM} \begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & 2^{\ell_P}\,\mathsf{poly}(T + \ell_P + 1/\Delta) \\ \mathsf{prover\ comm} & = & \mathsf{poly}(T + \ell_P + 1/\Delta) \\ \mathsf{verifier\ comm} & = & \mathsf{poly}(T + \ell_P + 1/\Delta) \\ \mathsf{verifier\ randomness} & = & \mathsf{poly}(T + \ell_P + 1/\Delta) \end{pmatrix}.$$

The resulting public-coin proof for $\overline{L}$ has completeness at least $2/3$ and soundness at most $1/3$. The same holds for promise problems $(L_{\mathrm{yes}}, L_{\mathrm{no}})$, with the resulting public-coin proof deciding $(L_{\mathrm{no}}, L_{\mathrm{yes}})$.

Proof of Theorem 5.1. Let $(\mathcal{P}, \mathcal{V})$ be the given one-round interactive proof for L, where $\mathcal{V} = (\mathcal{V}_1, \mathcal{V}_2)$. We first construct a protocol $(\mathcal{P}', \mathcal{V}')$ showing that

$$\overline{L} \in \mathsf{IP} \begin{pmatrix} \mathsf{rounds} & = & 2 \\ \mathsf{verifier\ time} & = & 2^{\ell_P}\,\mathsf{poly}(T + \ell_P + 1/\Delta) \\ \mathsf{prover\ comm} & = & \mathsf{poly}(T + \ell_P + 1/\Delta) \\ \mathsf{verifier\ comm} & = & \mathsf{poly}(T + \ell_P + 1/\Delta) \\ \mathsf{verifier\ randomness} & = & \mathsf{poly}(T + \ell_P + 1/\Delta) \\ \mathsf{completeness} & \ge & 1 - \epsilon'_c \\ \mathsf{soundness} & \le & \epsilon'_s \end{pmatrix},$$

where

$$1 - \epsilon'_c - \epsilon'_s \ge \Omega\left(\frac{\Delta^2}{\ell_P + \log(1/\Delta)}\right).$$

We then amplify this protocol and apply the private-to-public-coin transformation and round collapse.

Fix an input $x \in \{0,1\}^n$. Let ℓ_V be the length of the verifier's first message pk, and let r be the length of its random tape ρ; ℓ_P is the length of the prover's reply π. We have $\ell_V, r = O(T)$. The tape ρ includes all randomness of $\mathcal{V}$, and $\mathcal{V}_2(x,\rho,\pi) = 1$ denotes acceptance.[13] For each $\mathsf{pk} \in \{0,1\}^{\ell_V}$, define

$$S_{\mathsf{pk}} = \big\{\rho \in \{0,1\}^r : \mathcal{V}_1(x;\rho) = \mathsf{pk}\big\}.$$

Whenever $S_{\mathsf{pk}} \neq \varnothing$, for each $\pi \in \{0,1\}^{\ell_P}$ define

$$p_{\mathsf{pk},\pi} = \Pr_{\rho \leftarrow S_{\mathsf{pk}}}\big[\mathcal{V}_2(x,\rho,\pi) = 0\big] = \frac{1}{|S_{\mathsf{pk}}|}\sum_{\rho \in S_{\mathsf{pk}}}\big(1 - \mathcal{V}_2(x,\rho,\pi)\big).$$

We use the midpoint

$$\alpha = \frac{\epsilon_c + 1 - \epsilon_s}{2}$$

as the acceptance threshold. The protocol uses parameters

$$\begin{aligned} &\eta = \Theta(\Delta), && d = \Theta\big(\ell_P + \log(1/\Delta)\big), \\ &B = \Theta\big(d/\Delta^2\big), && N = \Theta\big(\Delta^{-2}\log B\big), \end{aligned} \tag{1}$$

The implicit constant in η is sufficiently small, so that $\Delta/2 - \eta = \Theta(\Delta)$. The constants in d, B, N are sufficiently large.

We use the hashing and concentration bounds from Lemmas 4.36, 4.33, and 4.34.

1. The verifier $\mathcal{V}'$ samples independent tapes $\rho_1, \dots, \rho_N \leftarrow \{0,1\}^r$ and sends $(\mathsf{pk}_i)_{i\in[N]}$, where $\mathsf{pk}_i = \mathcal{V}_1(x;\rho_i)$. It keeps the tapes ρ_i private.

2. The prover $\mathcal{P}'$ sends $(K_i)_{i\in[N]}$, where $K_i = |S_{\mathsf{pk}_i}|$.

3. The verifier $\mathcal{V}'$ checks that $1 \le K_i \le 2^r$ for every i, rejecting otherwise, and chooses $k_i \ge 0$ such that the expected bucket size $K_i 2^{-k_i}$ is at most a sufficiently small constant multiple of B, and is $\Theta(B)$ if $k_i > 0$. Independently for each i, it samples a hash

   $$h_i : \{0,1\}^r \to \{0,1\}^{k_i}$$

   from a $(2d+1)$-wise independent family as in Lemma 4.36. It sets $t_i = h_i(\rho_i)$ and sends the seeds specifying the hashes and the targets $(t_i)_{i\in[N]}$.

4. The prover $\mathcal{P}'$ computes the buckets

   $$C_i = \big\{\rho \in S_{\mathsf{pk}_i} : h_i(\rho) = t_i\big\} \qquad (i \in [N]).$$

   If any $|C_i| > B$, it aborts. Otherwise, it sends lists $(L_i)_{i\in[N]}$, with L_i containing every element of C_i exactly once.

5. Upon receiving the lists, the verifier $\mathcal{V}'$ does the following:

   - It checks that each L_i contains between 1 and B distinct r-bit strings and that

     $$\mathcal{V}_1(x;\rho) = \mathsf{pk}_i, \quad h_i(\rho) = t_i \quad \text{for every } \rho \in L_i, \qquad \rho_i \in L_i.$$

[13] Without loss of generality, the verifier's private state is $\mathsf{vk} = \rho$.

It rejects if any check fails.

- For every $i \in [N]$ and $\pi \in \{0,1\}^{\ell_P}$, it computes

$$\widehat{p}_{i,\pi} = \frac{1}{|L_i|} \sum_{\rho \in L_i} \big(1 - \mathcal{V}_2(x, \rho, \pi)\big), \qquad \widehat{q}_i = \min_{\pi \in \{0,1\}^{\ell_P}} \widehat{p}_{i,\pi}.$$

- It accepts if and only if

$$\frac{1}{N} \sum_{i=1}^{N} \widehat{q}_i \geq \alpha.$$

Efficiency. We first analyze the efficiency of the protocol.

- The prover clearly sends $O(NBr)$ bits.
- The verifier's running time is $2^{\ell_P}\,\mathsf{poly}(T + \ell_P + N + d + B)$: it enumerates all 2^{ℓ_P} replies and evaluates $\mathcal{V}_2$ on at most NB tapes for each reply. The list checks, hash evaluations, and arithmetic take polynomial time per reply or per transcript.
- By the hash-seed bound, the verifier sends $O(N(\ell_V + dr))$ bits.
- Sampling the private tapes and hash seeds uses $O(Ndr)$ verifier random bits, by the same bound.

Completeness. Suppose $x \notin L$ and $\mathcal{P}'$ is honest. If $k_i = 0$, then $C_i = S_{\mathsf{pk}_i}$, so bucket i has size $K_i < B$ and gives exact estimates $\widehat{p}_{i,\pi} = p_{\mathsf{pk}_i,\pi}$; we therefore consider only i with $k_i > 0$ below. It is easy to see that, if no abort occurs, all consistency checks pass. It therefore suffices to bound the probability that either of the following occurs:

(a) $|C_i| > B$ for some $i \in [N]$, causing an abort in Step 4.

(b) No abort occurs, but the final test in Step 5 fails, namely

$$\frac{1}{N} \sum_{i=1}^{N} \widehat{q}_i < \alpha.$$

Case (a). Fix $i \in [N]$ and ρ_i, and condition on $h_i(\rho_i) = t_i$. The tape ρ_i always belongs to C_i, and each other tape in S_{pk_i} belongs to C_i with probability 2^{-k_i}. Hence,

$$\mathbb{E}[|C_i|] = 1 + (K_i - 1)2^{-k_i}.$$

By the choice of k_i in Step 3, this expectation is at most a sufficiently small constant multiple of B. Since h_i is $(2d+1)$-wise independent, the variables

$$\mathbf{1}\big[h_i(\rho) = t_i\big], \qquad \rho \in S_{\mathsf{pk}_i} \setminus \{\rho_i\},$$

are $2d$-wise independent under the conditioning. Their sum is $|C_i| - 1$. Applying Lemma 4.34 with $k = 2d$ and deviation $\tau = \Theta(B)$ gives

$$\Pr[|C_i| > B] \leq \left(O\left(\frac{d}{B} + \frac{d^2}{B^2}\right)\right)^d.$$

This bound holds for every choice of the conditioned tape and hash value. Taking a union bound over all buckets, we obtain

$$\Pr\big[\exists i \in [N] : |C_i| > B\big] \le N \left(O\!\left(\frac{d}{B} + \frac{d^2}{B^2} \right) \right)^d . \tag{2}$$

Case (b). We first bound the error in the rejection-rate estimates. Our parameter choice makes ηB larger than a sufficiently large absolute constant. Define

$$\widehat{p}_{i,\pi} = \frac{1}{|C_i|} \sum_{\rho \in C_i} \big(1 - \mathcal{V}_2(x, \rho, \pi)\big)$$

using the full buckets, even if the protocol aborts. These agree with the protocol's estimates whenever no abort occurs.

Fix $i \in [N]$ and ρ_i. We now assume that $h_i(\rho_i) = t_i$ for a fixed target t_i (all expectations and probabilities below are conditioned on this). Since $K_i = |S_{\mathsf{pk}_i}|$, the choice of k_i in Step 3 gives

$$|S_{\mathsf{pk}_i}| 2^{-k_i} = K_i 2^{-k_i} = \Theta(B).$$

The tape ρ_i always belongs to C_i, and each other tape in S_{pk_i} belongs to C_i with probability 2^{-k_i}. Hence, for every fixed π,

$$\begin{aligned} \mathbb{E}[|C_i|] &= 1 + \big(|S_{\mathsf{pk}_i}| - 1\big) 2^{-k_i} = |S_{\mathsf{pk}_i}| 2^{-k_i} + 1 - 2^{-k_i}, \\ \mathbb{E}[|C_i| \widehat{p}_{i,\pi}] &= |S_{\mathsf{pk}_i}| 2^{-k_i} p_{\mathsf{pk}_i,\pi} + (1 - 2^{-k_i})\big(1 - \mathcal{V}_2(x, \rho_i, \pi)\big). \end{aligned}$$

In particular, these expectations differ from $|S_{\mathsf{pk}_i}| 2^{-k_i}$ and $|S_{\mathsf{pk}_i}| 2^{-k_i} p_{\mathsf{pk}_i,\pi}$, respectively, by at most 1.

Since h_i is $(2d+1)$-wise independent, the random variables

$$\mathbf{1}\big[h_i(\rho) = t_i\big], \qquad \rho \in S_{\mathsf{pk}_i} \setminus \{\rho_i\},$$

are $2d$-wise independent after conditioning on $h_i(\rho_i) = t_i$. Both $|C_i| - 1$ and $|C_i| \widehat{p}_{i,\pi} - (1 - \mathcal{V}_2(x, \rho_i, \pi))$ are sums of these variables, with expectation at most $|S_{\mathsf{pk}_i}| 2^{-k_i}$. Apply Lemma 4.34 to each sum with $k = 2d$ and deviation $\tau = \Theta(\eta |S_{\mathsf{pk}_i}| 2^{-k_i})$, choosing the implicit constant sufficiently small. Since $\eta |S_{\mathsf{pk}_i}| 2^{-k_i}$ is sufficiently large, the unit biases above can be absorbed into this deviation. Each required concentration bound fails with probability at most

$$\left(O\!\left(\frac{d}{\eta^2 B} + \frac{d^2}{\eta^2 B^2} \right) \right)^d .$$

Outside these failures, the size and numerator differ from their respective targets by at most a sufficiently small constant multiple of $\eta |S_{\mathsf{pk}_i}| 2^{-k_i}$. Consequently,

$$\begin{aligned} |\widehat{p}_{i,\pi} - p_{\mathsf{pk}_i,\pi}| &\le \frac{\big| |C_i| \widehat{p}_{i,\pi} - |S_{\mathsf{pk}_i}| 2^{-k_i} p_{\mathsf{pk}_i,\pi} \big| + p_{\mathsf{pk}_i,\pi} \big| |C_i| - |S_{\mathsf{pk}_i}| 2^{-k_i} \big|}{|C_i|} \\ &\le \eta, \end{aligned}$$

using $p_{\mathsf{pk}_i,\pi} \le 1$ and $|C_i| = \Theta(|S_{\mathsf{pk}_i}| 2^{-k_i})$. The constant in the deviation threshold is chosen small enough for the last inequality. A union bound over the N size estimates and $N 2^{\ell_P}$ numerator estimates bounds the probability of any failure by $N(2^{\ell_P} + 1)$ times the bound above.

Define $q_i = \min_{\pi \in \{0,1\}^{\ell_P}} p_{\mathsf{pk}_i,\pi}$. By soundness of $(\mathcal{P}, \mathcal{V})$, we have $\mathbb{E}[q_i] \geq 1 - \epsilon_s$: a prover can choose a reply minimizing $p_{\mathsf{pk}_i,\pi}$ for each first message pk_i. The q_i are independent and take values in $[0,1]$. Since $\Delta/2 - \eta = \Theta(\Delta)$, Hoeffding's inequality gives

$$\Pr\left[\frac{1}{N}\sum_{i=1}^{N} q_i < \alpha + \eta\right] \leq e^{-\Omega(N\Delta^2)}.$$

Whenever no abort occurs and all the estimates satisfy $|\widehat{p}_{i,\pi} - p_{\mathsf{pk}_i,\pi}| \leq \eta$, we have $\widehat{q}_i \geq q_i - \eta$ for every i. If also $\frac{1}{N}\sum_i q_i \geq \alpha + \eta$, then $\frac{1}{N}\sum_i \widehat{q}_i \geq \alpha$, so the final test in Step 5 passes. Taking a union bound, we obtain

$$\Pr[\text{Case (b)}] \leq N(2^{\ell_P} + 1)\left(O\left(\frac{d}{\eta^2 B} + \frac{d^2}{\eta^2 B^2}\right)\right)^d + e^{-\Omega(N\Delta^2)}. \tag{3}$$

Combining (2) and (3), and absorbing the overflow term into the estimation bound, the completeness error of $(\mathcal{P}', \mathcal{V}')$ satisfies

$$\epsilon_c' \leq N(2^{\ell_P} + 1)\left(O\left(\frac{d}{\eta^2 B} + \frac{d^2}{\eta^2 B^2}\right)\right)^d + e^{-\Omega(N\Delta^2)}. \tag{4}$$

We now check the parameter choice in (1). Since $\eta = \Theta(\Delta)$ and $B = \Theta(d/\Delta^2)$, choosing the constant in B sufficiently large makes the base of the moment term a sufficiently small absolute constant. Thus,

$$\epsilon_c' \leq N\, 2^{\ell_P - \Omega(d)} + e^{-\Omega(N\Delta^2)} \leq B^{-\Omega(1)}. \tag{5}$$

The exponent in the last bound can be made arbitrarily large by increasing the constants in d and N. Indeed,

$$\log(NB) = O\big(\log d + \log(1/\Delta)\big),$$

so $d = \Theta(\ell_P + \log(1/\Delta))$ controls both the union bound over replies and the desired inverse-polynomial dependence on B. The choice $N = \Theta(\Delta^{-2}\log B)$ gives the same control over the Hoeffding term.

Soundness. Suppose $x \in L$, and fix an arbitrary prover strategy. For acceptance, each claim must satisfy $1 \leq K_i \leq 2^r$, and each list L_i must contain between 1 and B distinct elements of C_i. We consider the following two cases:

(a) $L_i = C_i$ for every $i \in [N]$.

(b) $L_i \subsetneq C_i$ for some $i \in [N]$.

Case (a). Define the full-bucket estimates

$$Q_i = \min_{\pi \in \{0,1\}^{\ell_P}} \frac{1}{|C_i|}\sum_{\rho \in C_i}\big(1 - \mathcal{V}_2(x,\rho,\pi)\big).$$

These are defined for every choice of the submitted lists, and $Q_i = \widehat{q}_i$ in Case (a). Fix the prover's coins; we may assume its claims pass the range check, since otherwise the verifier rejects.

Condition on the first messages $(\mathsf{pk}_j)_j$, claims $(K_j)_j$, and hashes $(h_j)_j$. The tapes ρ_i remain independent and uniform in S_{pk_i}. For each i, fix a reply π_i^* attaining $q_i = \min_\pi p_{\mathsf{pk}_i,\pi}$. Each nonempty

bucket C of h_i on S_{pk_i} is selected with probability $|C|/|S_{\mathsf{pk}_i}|$. Thus, with expectations conditioned as above,

$$\begin{aligned}\mathbb{E}[Q_i] &\le \mathbb{E}\left[\frac{1}{|C_i|}\sum_{\rho\in C_i}\big(1-\mathcal{V}_2(x,\rho,\pi_i^*)\big)\right]\\ &= \frac{1}{|S_{\mathsf{pk}_i}|}\sum_{\rho\in S_{\mathsf{pk}_i}}\big(1-\mathcal{V}_2(x,\rho,\pi_i^*)\big) = p_{\mathsf{pk}_i,\pi_i^*} = q_i.\end{aligned}$$

Under the same conditioning, the Q_i are independent and take values in $[0,1]$. Whenever $\frac{1}{N}\sum_i q_i \le \epsilon_c+\eta$, Hoeffding's inequality therefore bounds the probability that $\frac{1}{N}\sum_i Q_i \ge \alpha$ by $e^{-\Omega(N\Delta^2)}$.

Over the sampled first messages, the q_i are independent and, by completeness of $(\mathcal{P},\mathcal{V})$, satisfy $\mathbb{E}[q_i]\le\epsilon_c$. Another application of Hoeffding's inequality gives

$$\Pr\left[\frac{1}{N}\sum_i q_i > \epsilon_c+\eta\right] \le e^{-\Omega(N\Delta^2)}.$$

In Case (a), acceptance requires $\frac{1}{N}\sum_i Q_i \ge \alpha$. Taking a union bound, we obtain

$$\Pr[\mathcal{V}' \text{ accepts and Case (a) occurs}] \le e^{-\Omega(N\Delta^2)}. \tag{6}$$

Case (b). Fix the prover's coins and the visible transcript through Step 4, and choose an i with $L_i \subsetneq C_i$. Conditional on this transcript, ρ_i is uniform in C_i. Hence,

$$\Pr[\rho_i \notin L_i \mid \text{visible transcript}] = 1-\frac{|L_i|}{|C_i|} \ge \frac{1}{|L_i|+1} \ge \frac{1}{B+1}, \tag{7}$$

since $|C_i| \ge |L_i|+1$ and $|L_i| \le B$. Thus, in Case (b), the verifier rejects with probability at least $1/(B+1)$.

Combining (6) and (7), the soundness error of $(\mathcal{P}',\mathcal{V}')$ satisfies

$$\begin{aligned}\epsilon_s' &\le 1-\frac{1}{B+1}+e^{-\Omega(N\Delta^2)}\\ &\le 1-\frac{1}{B+1}+B^{-\Omega(1)}.\end{aligned} \tag{8}$$

As above, the exponent in the error term can be made arbitrarily large by increasing the constant in N.

Amplification and public coins. Choose the constants in d and N so that the sum of the completeness error in (5) and the additive soundness error term in (8) is a sufficiently small constant multiple of $1/(B+1)$. The resulting gap is

$$1-\epsilon_c'-\epsilon_s' \ge \Omega\left(\frac{1}{B+1}\right) = \Omega\left(\frac{\Delta^2}{d}\right).$$

Applying Lemma 4.29, followed by Theorems 4.30 and 4.31, gives the claimed one-round public-coin proof with the stated resource bounds. □

Theorem 5.2. *Suppose L has a one-round interactive proof with verifier time $T(n)\ge n$, prover-message length $\ell_{\mathcal{P}}(n)\ge 1$, and error bounds $\epsilon_c(n),\epsilon_s(n)\in[0,1]$ with $\Delta(n)=1-\epsilon_c(n)-\epsilon_s(n)>0$. If the*

first-message map depends on x only through its length, namely $\mathcal{V}_1(x;\rho) = \mathcal{V}_1(1^n;\rho)$ for every $x \in \{0,1\}^n$ and random tape ρ, then

$$L \in \mathsf{SIZE}\Big[2^{\ell_{\mathcal{P}}(n)}\,\mathsf{poly}\big(T(n) + \ell_{\mathcal{P}}(n) + 1/\Delta(n)\big)\Big].$$

Moreover, L can be decided deterministically in this time with advice of length $\mathsf{poly}(T+\ell_{\mathcal{P}}+1/\Delta)$ depending only on n.

Proof. Fix n and suppress it from the resource bounds. Let $\rho \in \{0,1\}^r$ contain all verifier randomness, where $r = O(T)$, and write $\mathcal{V}_2(x,\rho,\pi) = 1$ for acceptance. For each first message pk, let

$$S_{\mathsf{pk}} = \{\rho \in \{0,1\}^r : \mathcal{V}_1(1^n;\rho) = \mathsf{pk}\}.$$

For $S_{\mathsf{pk}} \neq \varnothing$, define

$$p_x(\mathsf{pk},\pi) = \Pr_{\rho \leftarrow S_{\mathsf{pk}}}[\mathcal{V}_2(x,\rho,\pi) = 1], \qquad v(x) = \mathbb{E}_{\mathsf{pk}}\left[\max_{\pi \in \{0,1\}^{\ell_{\mathcal{P}}}} p_x(\mathsf{pk},\pi)\right],$$

where pk has the distribution of an honestly generated first message. Since the prover can choose its reply separately for each pk, $v(x)$ is its optimal acceptance probability. Thus $v(x) \geq 1 - \epsilon_c$ for $x \in L$ and $v(x) \leq \epsilon_s$ for $x \notin L$.

Choose $\eta = \Theta(\Delta)$ with $\eta < \Delta/4$, and integers

$$N = \Theta\left(\frac{n}{\eta^2}\right), \qquad M = \Theta\left(\frac{n + \ell_{\mathcal{P}} + \log N}{\eta^2}\right).$$

Sample N independent first messages $\mathsf{pk}_1, \ldots, \mathsf{pk}_N$. For each i, independently sample M uniform tapes $\rho_{i,1}, \ldots, \rho_{i,M}$ from S_{pk_i}. These distributions depend only on n. Given the messages and tapes as advice, compute

$$\widehat{v}(x) = \frac{1}{N}\sum_{i=1}^{N} \max_{\pi \in \{0,1\}^{\ell_{\mathcal{P}}}} \frac{1}{M}\sum_{j=1}^{M} \mathcal{V}_2(x,\rho_{i,j},\pi).$$

By Lemma 4.33 and a union bound over all x, the average $N^{-1}\sum_i \max_\pi p_x(\mathsf{pk}_i,\pi)$ is within η of $v(x)$ simultaneously for every x, except with probability at most $2^{n+1}e^{-2N\eta^2}$. Conditioned on the messages, another union bound over all i, x, π shows that every inner average in $\widehat{v}(x)$ is within η of $p_x(\mathsf{pk}_i,\pi)$, except with probability at most $N2^{n+\ell_{\mathcal{P}}+1}e^{-2M\eta^2}$. Taking sufficiently large constants in N, M, the sum of these failure probabilities is less than 1. Hence there is a fixed choice of advice for which

$$|\widehat{v}(x) - v(x)| \leq 2\eta < \Delta/2 \qquad \text{for every } x \in \{0,1\}^n.$$

Accepting exactly when $\widehat{v}(x) > (1 - \epsilon_c + \epsilon_s)/2$ therefore decides L.

The advice has length $O(NMT)$; it need not be efficiently sampled. Computing $\widehat{v}(x)$ uses $NM2^{\ell_{\mathcal{P}}}$ verifier evaluations. Since $T \geq n$, the advice length and running time satisfy the claimed bounds. Hardwiring the advice and implementing each verifier evaluation by a circuit of size $\mathsf{poly}(T)$ gives the circuit-size bound. □

6 From One-Query Reductions to Laconic Interactive Proofs

Throughout this section, $\ell_{\mathcal{P}}(\lambda, n)$ bounds the prover-message length over all witnesses of length at most $q(n)$. Succinctness gives $\ell_{\mathcal{P}}(\lambda, n) = \mathsf{poly}(\lambda)(n + q(n))^{o(1)}$.

We require that whenever a prover has nonnegligible soundness advantage over c_0, the reduction achieves nonnegligible advantage against the target game. Both are measured in the source security parameter λ. Each reduction runs the target challenger at $\lambda' = f(1^\lambda, n)$ for its target-parameter map f, and the hardness hypothesis is security against the honest-prover reduction (Definition 6.1), which is implied by T-hardness at f (Definition 4.12).

Security against the honest-prover reduction. Our lower bounds use the security of $(\mathcal{C}, c)$ only against the reduction run with the honest prover.

Definition 6.1 (Security against the honest-prover reduction)**.** *Let $\Pi = (\mathcal{V}, \mathcal{P})$ be a two-message laconic argument with soundness baseline c_0 for an* NP *relation $\mathcal{R}$ with witness bound q, and let $(\mathcal{R}, F)$ be a k-query, T-time black-box reduction from its soundness to $(\mathcal{C}, c)$ with target-parameter map f. For a statement x and a string w, let $\widetilde{\mathcal{P}}_{x,w}$ answer each query pk with $\mathcal{P}(\mathsf{pk}, x, w)$ if $(x, w) \in \mathcal{R}$, and with a fixed dummy response otherwise.*

The honest-prover reduction $\mathcal{R} \circ \mathcal{P}$ *takes auxiliary input $(x, z, w_1, \ldots, w_k)$, runs $\mathcal{R}(1^\lambda, x, z)$, and answers its jth query using $\widetilde{\mathcal{P}}_{x,w_j}$. Its allowed auxiliary inputs at (λ, n) are those with $|x| = n$, $|z| \le T(\lambda, n)$, and $|w_j| \le q(n)$. We say that $(\mathcal{C}, c)$ is* secure against $\mathcal{R} \circ \mathcal{P}$ at f *if it is T-secure against this algorithm (Definition 4.14). More generally, for a bound $\Gamma(\lambda, n) \ge \lambda$, it is* Γ-secure against $\mathcal{R} \circ \mathcal{P}$ at f *if it is Γ-secure against this algorithm.*

For adaptive soundness, x is replaced by 1^n and each w_j by a pair (x_j, w_j) with $|x_j| = n$, and the jth query is answered with $(x_j, \widetilde{\mathcal{P}}_{x_j,w_j}(\mathsf{pk}))$.

For $x \in L$ with valid witnesses, $\mathcal{R} \circ \mathcal{P}$ is the reduction run with honest provers, which we use for completeness. Security against $\mathcal{R} \circ \mathcal{P}$ is also the hypothesis of the following lemma. Since $\mathcal{R} \circ \mathcal{P}$ is computed by circuits of size $\mathsf{poly}(T(\lambda, n))$, T-hardness at f (Definition 4.12) implies security against $\mathcal{R} \circ \mathcal{P}$ at f.

Lemma 6.2. *Let $\Pi = (\mathcal{V}, \mathcal{P})$ be a two-message laconic argument with soundness baseline c_0 for an* NP *relation $\mathcal{R}$ with witness bound q. Let $\Gamma(\lambda, n) \ge \lambda$, and let $(\mathcal{R}, F)$ be a non-uniform, k-query, T-time black-box reduction from the non-adaptive or adaptive soundness of Π to the Γ-hardness of $(\mathcal{C}, c)$ at its target-parameter map f (Definition 4.13), with possibly randomized advice F. Let X be the corresponding index set: the false statements for non-adaptive soundness, and $\{1^n : n \in \mathbb{N}\}$ for adaptive soundness. Fix a nondecreasing polynomial p and a constant $\eta \in (0, 1]$, and assume that $(\mathcal{C}, c)$ is Γ-secure against $\mathcal{R} \circ \mathcal{P}$ at f.*

Then there exists a constant C such that, for all sufficiently large λ, every $x \in X$ with $|x| \le p(\lambda)$, and every possibly unbounded randomized response oracle $\mathcal{P}^*$ with soundness advantage at least η at (λ, x),

$$\Pr\left[\mathcal{R}^{\mathcal{P}^*}(1^\lambda, x, F(1^\lambda, x)) \text{ wins } \mathcal{C}(1^{f(1^\lambda, |x|)})\right] \ge c + \Gamma(\lambda, |x|)^{-C}.$$

The constant C and the cutoff may depend on p, η, Γ, and the fixed reduction, but are independent of x and $\mathcal{P}^*$.

For a reduction to $(\mathcal{C}, c)$ (Definition 4.13), we apply the lemma with $\Gamma(\lambda, n) = \lambda$. Its hypothesis then follows from security against $\mathcal{R} \circ \mathcal{P}$, since $T \ge \lambda$, and its conclusion is an inverse-polynomial advantage bound λ^{-C}.

Proof. The proof is a diagonalization argument in the spirit of [Bel02]. Suppose otherwise. Then for every integer $i \geq 1$ there are infinitely many λ with a counterexample for $C = i$. Choose increasing parameters $2 \leq \lambda_1 < \lambda_2 < \cdots$, indices $x_i \in X$ of length at most $p(\lambda_i)$, and response oracles $\mathcal{P}_i^*$ with soundness advantage at least η, but for which the reduction wins with probability less than $c + \Gamma(\lambda_i, |x_i|)^{-i}$.

Define a non-uniform adversary family using $\mathcal{P}_i^*$ at parameters λ_i, and the dummy oracle that answers every query with the fixed dummy response of Definition 6.1 elsewhere. Use indices x_i at the selected parameters and x_1 elsewhere. Since p is nondecreasing, these indices have length at most $p(\lambda)$ for every $\lambda \geq \lambda_1$, so the index sequence is polynomially bounded. The soundness advantage is at least η at every λ_i, so it is nonnegligible. Since $(\mathcal{R}, F)$ is a reduction to the Γ-hardness of $(\mathcal{C}, c)$, it Γ-breaks $(\mathcal{C}, c)$: there is a polynomial p', and hence a constant C with $p'(t) \leq t^C$ for all $t \geq 2$, such that $\delta'(\lambda) > \Gamma(\lambda, |x_\lambda|)^{-C}$ for infinitely many λ.

We show that $\delta'(\lambda) < \Gamma(\lambda, |x_\lambda|)^{-C}$ for all sufficiently large λ, a contradiction. At a selected parameter λ_i with $i > C$, we have $\delta'(\lambda_i) < \Gamma^{-i} \leq \Gamma^{-C}$, since $\Gamma \geq \lambda \geq 2$. At the remaining parameters $\lambda \geq \lambda_1$, the reduction runs with the dummy oracle. This is exactly $\mathcal{R} \circ \mathcal{P}$ on auxiliary input $(x_1, z, w_1, \ldots, w_k)$ with invalid witnesses w_j, for instance w_j empty; for adaptive soundness, take invalid pairs (x_j, w_j). Its winning probability, averaged over $z \leftarrow F(1^\lambda, x_1)$, is at most its maximum over advice strings with $|z| \leq T$, which is at most $c + \Gamma^{-(C+1)} < c + \Gamma^{-C}$ for all sufficiently large λ by Γ-security against $\mathcal{R} \circ \mathcal{P}$, applied to the worst allowed advice at each λ. □

6.1 Non-Uniform Reductions for Non-Adaptive Soundness

We consider a non-uniform reduction whose advice has description complexity s_F, as in Definition 4.15. Here T bounds the running time of the original reduction, and we account for the sampler description separately. The mixed verifier takes this description as an explicit input and samples the advice privately. We show completeness for every candidate description and statistical soundness for a description of the reduction's advice sampler.

Theorem 6.3. *Fix a nondecreasing polynomial p and a security-parameter function $\lambda = \lambda(n)$ such that $n \leq p(\lambda(n))$ for every n. Suppose L has a two-message laconic argument $\Pi = (\mathcal{V} = (\mathcal{V}_1, \mathcal{V}_2), \mathcal{P})$ with soundness baseline c_0, verifier-message length $\ell_{\mathcal{V}}(\lambda, n)$ and prover-message length $\ell_{\mathcal{P}}(\lambda, n)$, where $n = |x|$. Suppose its non-adaptive soundness admits a non-uniform, one-query, $T(\lambda, n)$-time black-box reduction $(\mathcal{R}, F)$ with advice of description complexity $s_F(\lambda, n)$ to the Γ-hardness of a game-based assumption $(\mathcal{C}, c)$ at its target-parameter map f (Definition 4.13), for a bound $\lambda \leq \Gamma(\lambda, n) \leq T(\lambda, n)$. Assume that the completeness error κ of Π satisfies, for every constant $K > 0$, $\kappa(\lambda, n) \leq \Gamma(\lambda, n)^{-K}$ for all sufficiently large λ and all $n \leq p(\lambda)$.*

For each $x \notin L$, fix a circuit description $D_{\lambda,x}$ of the sampler $\rho \mapsto F(1^\lambda, x; \rho)$, of length at most $s_F(\lambda, n)$. Define the promise problem

$$L^F_{\text{yes}} = \{(1^\lambda, x, D) : x \in L,\ |D| \leq s_F(\lambda, n)\},$$
$$L^F_{\text{no}} = \{(1^\lambda, x, D_{\lambda,x}) : x \notin L\},$$

where $\lambda = \lambda(|x|)$ in both sets. Then, assuming that $(\mathcal{C}, c)$ is secure against $\mathcal{R} \circ \mathcal{P}$ at the reduction's target-parameter map f (Definition 6.1), there exist completeness and soundness error bounds ϵ_c, ϵ_s satisfying

$$1 - \epsilon_c(n) - \epsilon_s(n) \geq \frac{1}{\mathsf{poly}(\Gamma(\lambda, n))}$$

such that

$$(L^F_{\text{yes}}, L^F_{\text{no}}) \in \mathsf{IP}\begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & \mathsf{poly}(T(\lambda,n), s_F(\lambda,n)) \\ \mathsf{prover\ comm} & = & \ell_{\mathcal{P}}(\lambda,n) \\ \mathsf{verifier\ comm} & = & \ell_{\mathcal{V}}(\lambda,n) \\ \mathsf{completeness} & \geq & 1-\epsilon_c \\ \mathsf{soundness} & \leq & \epsilon_s \end{pmatrix}.$$

The verifier is uniform: it receives D as part of the common input and uses it to sample the advice. The resource bounds are expressed in terms of the statement length $n = |x|$.

For a uniform reduction, the advice is empty, so $D_{\lambda,x}$ can be taken to be a fixed circuit that the verifier supplies itself. In this case, the theorem directly gives

$$L \in \mathsf{IP}\begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & \mathsf{poly}(T(\lambda,n)) \\ \mathsf{prover\ comm} & = & \ell_{\mathcal{P}}(\lambda,n) \\ \mathsf{verifier\ comm} & = & \ell_{\mathcal{V}}(\lambda,n) \\ \mathsf{completeness} & \geq & 1-\epsilon_c \\ \mathsf{soundness} & \leq & \epsilon_s \end{pmatrix}.$$

Proof. Let $\Pi = (\mathcal{V} = (\mathcal{V}_1, \mathcal{V}_2), \mathcal{P})$ be the given laconic argument for a fixed NP relation with witness bound q. Since $D_{\lambda,x}(\rho)$ is distributed as $F(1^\lambda, x)$, Lemma 6.2 applies to $(\mathcal{R}, F)$, and its hypothesis follows from security against $\mathcal{R} \circ \mathcal{P}$ (see the discussion after Definition 6.1). Apply it with $\eta = (1 - c_0)/2$. Its hypothesis, Γ-security against $\mathcal{R} \circ \mathcal{P}$, follows from security against $\mathcal{R} \circ \mathcal{P}$, since $\Gamma \leq T$. Write $\delta(\lambda, n) = \Gamma(\lambda, n)^{-C}$ for the resulting advantage bound, with C a constant, and choose a cutoff λ_0 so that the lemma applies for every $\lambda \geq \lambda_0$ and every false statement of length $n \leq p(\lambda)$.

Without loss of generality, $\mathcal{R}$ never queries strings longer than $\ell_{\mathcal{V}}(\lambda, n)$.

Fix a constant $\beta \in (0, 1)$ satisfying $\beta > 2c/(2c + 1 - c_0)$ and define $\tau = 1 - c(1 - \beta)$. We now construct a one-round laconic *proof* system $\Pi' = (\mathcal{V}' = (\mathcal{V}'_1, \mathcal{V}'_2), \mathcal{P}')$ for $(L^F_{\text{yes}}, L^F_{\text{no}})$ as follows.

On common input $(1^\lambda, x, D)$, the verifier rejects if $|D| > s_F(\lambda, n)$, and rejects prover messages longer than $\ell_{\mathcal{P}}(\lambda, n)$.

- The verifier $\mathcal{V}'_1(1^\lambda, x, D) \to (\mathsf{pk}, \mathsf{vk}')$ first does the following: Sample $b \leftarrow \{0, 1\}$, where $b = 0$ with probability β, and $b = 1$ with probability $1 - \beta$.
  - If $b = 0$, sample $(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x)$ and set $\mathsf{vk}' = (0, \mathsf{vk})$.
  - If $b = 1$, sample fresh coins ρ and set $z = D(\rho)$, interpreting malformed descriptions and outputs longer than $T(\lambda, n)$ as the empty string. Sample random tapes for $\mathcal{R}$ and $\mathcal{C}$, and simulate the interaction of $\mathcal{R}(1^\lambda, x, z)$ with $\mathcal{C}(1^{\lambda'})$, where $\lambda' = f(1^\lambda, n)$, until $\mathcal{R}$ makes its oracle query pk. Set $\mathsf{vk}' = (1, \mathsf{st}_{\mathcal{R}}, \mathsf{st}_{\mathcal{C}})$, where the saved states include both machines' random tapes. If the interaction terminates without an oracle query, set pk to a fixed dummy message of length $\ell_{\mathcal{V}}(\lambda, n)$ and record the terminal states and outcome in vk'.

  Send pk to the prover and keep vk' private.
- Upon receiving pk, the prover $\mathcal{P}'(\mathsf{pk}, x, w)$ outputs $\pi \leftarrow \mathcal{P}(\mathsf{pk}, x, w)$, ignoring D.

- Upon receiving π, the verifier $\mathcal{V}_2'(\mathsf{vk}', x, \pi)$ reads the branch bit b from vk' and does the following:
  - If $b = 0$, recover vk from vk' and output $\mathcal{V}_2(\mathsf{vk}, x, \pi)$.
  - If $b = 1$, supply π as the oracle answer and resume the simulation from the states and random tapes saved in vk', or ignore π if no oracle query was made. Reject if $\mathcal{C}$ outputs win in this simulation; otherwise, accept.

Completeness. We show that completeness must hold for *every* allowed D. Fix $x \in L$ of length $n \le p(\lambda)$, an NP witness w of length at most $q(n)$, and D with $|D| \le s_F(\lambda, n)$. By completeness of Π, when $b = 0$, $\mathcal{V}_2$ accepts with probability at least $1 - \kappa(\lambda, n)$, where $\kappa(\lambda, n) \le \Gamma(\lambda, n)^{-(C+1)}$ for all sufficiently large λ, uniformly over all allowed x, w for the fixed polynomial p.

When $b = 1$, fix a sampled advice string z with $|z| \le T(\lambda, n)$. The adversary that runs $\mathcal{R}(1^\lambda, x, z)$ and answers its oracle query using $\mathcal{P}(\mathsf{pk}, x, w)$ is $(\mathcal{R} \circ \mathcal{P})(x, z, w)$ (Definition 6.1), so security against $\mathcal{R} \circ \mathcal{P}$ (Definition 4.14) gives, for all sufficiently large λ,

$$\Pr\big[(\mathcal{R} \circ \mathcal{P})(x, z, w) \text{ wins } \mathcal{C}(1^{\lambda'})\big] \le c + \Gamma(\lambda, n)^{-(C+1)}.$$

Averaging over $z = D(\rho)$ preserves this bound for every D. The mixed verifier therefore accepts the honest prover with probability at least

$$\beta\big(1 - \Gamma^{-(C+1)}\big) + (1 - \beta)\big(1 - c - \Gamma^{-(C+1)}\big) \ge \tau - \Gamma(\lambda, n)^{-(C+1)}.$$

Statistical soundness. Fix $\lambda \ge \lambda_0$, a false statement $x \notin L$ of length $n \le p(\lambda)$, the description $D = D_{\lambda,x}$, and an arbitrary, possibly unbounded prover $\mathcal{P}^*$. The prover may know D, but with x, D fixed, its strategy is still an oracle to which Lemma 6.2 applies. Let

$$\alpha_{\text{cheat}} = \Pr\left[\ \mathcal{V}_2(\mathsf{vk}, x, \pi) = 1 \ \ : \ \begin{array}{c}(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x) \\ \pi \leftarrow \mathcal{P}^*(\mathsf{pk})\end{array}\ \right].$$

Since $D(\rho)$ has exactly the distribution $F(1^\lambda, x)$, the acceptance probability in Π' is exactly

$$\beta\alpha_{\text{cheat}} + (1 - \beta)\left(1 - \Pr\left[\mathcal{R}^{\mathcal{P}^*}(1^\lambda, x, z) \text{ wins } \mathcal{C}(1^{\lambda'})\right]\right), \tag{9}$$

where $z \leftarrow F(1^\lambda, x)$. We consider two cases, according to the success probability α_{cheat} on an honestly generated first message.

If $\alpha_{\text{cheat}} < (1 + c_0)/2$, the total acceptance probability is at most

$$\frac{\beta(1 + c_0)}{2} + (1 - \beta) = \tau - \left(\frac{\beta(1 - c_0)}{2} - (1 - \beta)c\right) < \tau$$

by our choice of $\beta > 2c/(1 - c_0 + 2c)$. Otherwise, $\alpha_{\text{cheat}} \ge (1 + c_0)/2$, so its soundness advantage is at least $(1 - c_0)/2$, and Lemma 6.2 gives

$$\Pr\left[\mathcal{R}^{\mathcal{P}^*}(1^\lambda, x, z) \text{ wins } \mathcal{C}(1^{\lambda'})\right] \ge c + \delta(\lambda, n).$$

Substituting into (9) and using $\alpha_{\text{cheat}} \le 1$, we obtain

$$\Pr[\text{accept}] \le \beta + (1 - \beta)\big(1 - c - \delta(\lambda, n)\big) = \tau - (1 - \beta)\delta(\lambda, n).$$

Both cases give acceptance probability at most $\tau - \xi(\lambda, n)$, where

$$\xi(\lambda, n) = \min\left\{\frac{\beta(1-c_0)}{2} - (1-\beta)c,\ (1-\beta)\delta(\lambda, n)\right\}.$$

The same description $D_{\lambda,x}$ works for all prover strategies.

Error bounds. Since $\xi(\lambda, n) \geq \min\{\Omega(1), (1-\beta)\Gamma(\lambda, n)^{-C}\}$ and $\Gamma \geq \lambda$, we can increase λ_0 so that $\Gamma(\lambda, n)^{-(C+1)} \leq \xi(\lambda, n)/2$ for every $\lambda \geq \lambda_0$ and every $n \leq p(\lambda)$. Taking

$$\epsilon_c(n) = 1 - \tau + \frac{\xi(\lambda(n), n)}{2}, \qquad \epsilon_s(n) = \tau - \xi(\lambda(n), n),$$

we obtain a completeness–soundness gap of at least $\xi(\lambda, n)/2 \geq 1/\operatorname{poly}(\Gamma(\lambda, n))$ whenever $\lambda(n) \geq \lambda_0$.

On the finitely many input lengths with $\lambda(n) < \lambda_0$, the verifier decides membership directly. Evaluating the sampler, computing $\lambda(n)$, running the mixed protocol, and takes time $\operatorname{poly}(T(\lambda, n), s_F(\lambda, n))$, and the prover and verifier message lengths remain $\ell_{\mathcal{P}}(\lambda, n)$ and $\ell_{\mathcal{V}}(\lambda, n)$. □

Combining with Theorem 5.1 and the round-collapse theorem, we obtain the following corollary.

Corollary 6.4. *Under the hypotheses of Theorem 6.3, we have*

$$L \in \mathsf{coAM}\begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & 2^{\ell_{\mathcal{P}}(\lambda,n)} \cdot \operatorname{poly}(T(\lambda, n), s_F(\lambda, n)) \\ \mathsf{prover\ comm} & = & \operatorname{poly}(T(\lambda, n), s_F(\lambda, n)) \\ \mathsf{verifier\ comm} & = & \operatorname{poly}(T(\lambda, n), s_F(\lambda, n)) \\ \mathsf{verifier\ randomness} & = & \operatorname{poly}(T(\lambda, n), s_F(\lambda, n)) \end{pmatrix}.$$

Proof. Write $n = |x|$ and $\lambda = \lambda(n)$. The proof of Theorem 5.1 applies to promise problems as well: it uses only completeness on yes-instances and soundness on no-instances. Applying it to the protocol of Theorem 6.3, gives a uniform one-round public-coin proof for $(L^F_{\mathrm{no}}, L^F_{\mathrm{yes}})$, with completeness at least $2/3$ and soundness at most $1/3$. Its verifier time is $2^{\ell_{\mathcal{P}}(\lambda,n)} \operatorname{poly}(T(\lambda, n), s_F(\lambda, n))$, and its communication and randomness are bounded by $\operatorname{poly}(T(\lambda, n), s_F(\lambda, n))$. Call its verifier $\mathcal{V}'$.

To obtain a protocol for $\overline{L}$ on common input x, the prover first sends a circuit description D of the sampler $\rho \mapsto F(1^\lambda, x; \rho)$. The verifier rejects if $|D| > s_F(\lambda, n)$; otherwise, the parties run the public-coin proof with verifier $\mathcal{V}'$ on common input $(1^\lambda, x, D)$.

If $x \notin L$, the honest prover sends $D = D_{\lambda,x}$, so $(1^\lambda, x, D) \in L^F_{\mathrm{no}}$ and completeness is at least $2/3$. If $x \in L$, every allowed D gives an input in L^F_{yes}, on which $\mathcal{V}'$ accepts every prover with probability at most $1/3$. Since D is sent before the public coins of $\mathcal{V}'$ are sampled, this soundness bound holds for every choice of the first message, including a randomized choice. Thus we have a three-message public-coin protocol for $\overline{L}$.

This yields a two-round public-coin proof (with an empty verifier message in the first round). Communicating D adds at most $s_F(\lambda, n)$ prover bits, so the total prover communication is $\operatorname{poly}(T(\lambda, n), s_F(\lambda, n))$. Applying Theorem 4.31 with $k = 2$ gives a one-round proof with verifier time $2^{\ell_{\mathcal{P}}(\lambda,n)} \operatorname{poly}(T(\lambda, n), s_F(\lambda, n))$ and the claimed prover communication. The round-collapse transformation also has only polynomial overhead in verifier communication and randomness [GVW02, Appendix B.3], giving the remaining claimed bounds. □

Corollary 6.5. *Under the hypotheses of Corollary 6.4, if F is deterministic, then*

$$L \in \mathsf{coAM}\begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & 2^{\ell_{\mathcal{P}}(\lambda,n)} \cdot \mathsf{poly}(T(\lambda,n)) \\ \mathsf{prover\ comm} & = & \mathsf{poly}(T(\lambda,n)) \\ \mathsf{verifier\ comm} & = & \mathsf{poly}(T(\lambda,n)) \\ \mathsf{verifier\ randomness} & = & \mathsf{poly}(T(\lambda,n)) \end{pmatrix}.$$

Proof. For each (λ, x), hardwire the string $F(1^\lambda, x)$, of length at most $T(\lambda, |x|)$, into a constant-output circuit. This gives description complexity $s_F(\lambda, n) = \mathsf{poly}(T(\lambda, n))$. Apply Corollary 6.4. □

6.2 Arbitrary Advice Distributions

We next use Carathéodory's theorem to approximate arbitrary advice distributions by samplers of bounded description complexity.

Theorem 6.6. *Fix λ, x and a one-query T-time reduction $\mathcal{R}$ interacting with a challenger $\mathcal{C}$, with queries in $\{0,1\}^{\ell_{\mathcal{V}}}$ and replies in $\{0,1\}^{\ell_{\mathcal{P}}}$, where $\ell_{\mathcal{V}}, \ell_{\mathcal{P}} \le T$. Let $\mathcal{D}_{x,\lambda}$ be any distribution on advice strings of length at most T, sampled independently of all other coins. For every $\varepsilon \in (0,1)$, there is a distribution $\widetilde{\mathcal{D}}_{x,\lambda}$ generated by a sampler of description complexity*

$$2^{\ell_{\mathcal{P}}+\ell_{\mathcal{V}}}\,\mathsf{poly}\big(T, \log(1/\varepsilon)\big),$$

such that replacing $z \leftarrow \mathcal{D}_{x,\lambda}$ by $z \leftarrow \widetilde{\mathcal{D}}_{x,\lambda}$ changes the winning probability of $\mathcal{R}^{\mathcal{A}}(1^\lambda, x, z)$ against $\mathcal{C}(1^{f(1^\lambda,|x|)})$ by at most ε for every possibly unbounded randomized response oracle $\mathcal{A}$. The replacement distribution is independent of $\mathcal{A}$.

Proof. Suppress x, λ and put $d = 2^{\ell_{\mathcal{P}}+\ell_{\mathcal{V}}}$. For the analysis, regard an execution with no oracle query as making the dummy query $0^{\ell_{\mathcal{V}}}$ and ignoring its reply, keeping the recorded outcome. For each advice string z, define $v_z \in [0,1]^d$ by

$$v_z(\mathsf{pk}, \pi) = \Pr\big[\text{the query is } \mathsf{pk} \text{ and } \mathcal{C} \text{ outputs } \mathsf{win} \text{ when the reply is } \pi\big].$$

The reply is fixed to π, and the probability is over the reduction and challenger coins. Since the query is chosen before the oracle replies, the winning probability with advice distribution $\mathcal{D}$ is

$$\sum_{\mathsf{pk},\pi} \Pr[\mathcal{A}(\mathsf{pk}) = \pi]\, \mathbb{E}_{z \leftarrow \mathcal{D}}[v_z(\mathsf{pk}, \pi)].$$

Thus it suffices to preserve $\mathbb{E}_{z\leftarrow\mathcal{D}}[v_z]$. This vector lies in the convex hull of $\{v_z : z \in \mathrm{supp}(\mathcal{D})\}$. By Carathéodory's theorem [COA20, Theorem 2], there are $m \le d+1$ advice strings $z_1, \ldots, z_m$ and weights $\alpha_1, \ldots, \alpha_m \ge 0$ summing to one such that

$$\mathbb{E}_{z\leftarrow\mathcal{D}}[v_z] = \sum_{i=1}^{m} \alpha_i v_{z_i}.$$

Let $\widehat{\mathcal{D}}$ assign weight α_i to z_i. The same equality preserves the winning probability for every $\mathcal{A}$.

To obtain a circuit sampler, apply Lemma 4.37 to $\widehat{\mathcal{D}}$, with strings of length at most T. This gives a distribution $\widetilde{\mathcal{D}}$ at statistical distance at most ε from $\widehat{\mathcal{D}}$, so every winning probability changes by at most ε. Since $m \le d+1$, its sampler has description complexity $2^{\ell_{\mathcal{P}}+\ell_{\mathcal{V}}}\,\mathsf{poly}(T, \log(1/\varepsilon))$. □

Corollary 6.7. *Under the hypotheses of Corollary 6.4, but allowing an arbitrary randomized advice mapping F without a bound on its description complexity, we have*

$$L \in \mathsf{coAM}\begin{pmatrix}\mathsf{rounds} & = & 1\\ \mathsf{verifier\ time} & = & 2^{O(\ell_{\mathcal{P}}(\lambda,n)+\ell_{\mathcal{V}}(\lambda,n))}\cdot \mathsf{poly}(T(\lambda,n))\\ \mathsf{prover\ comm} & = & 2^{O(\ell_{\mathcal{P}}(\lambda,n)+\ell_{\mathcal{V}}(\lambda,n))}\cdot \mathsf{poly}(T(\lambda,n))\\ \mathsf{verifier\ comm} & = & 2^{O(\ell_{\mathcal{P}}(\lambda,n)+\ell_{\mathcal{V}}(\lambda,n))}\cdot \mathsf{poly}(T(\lambda,n))\\ \mathsf{verifier\ randomness} & = & 2^{O(\ell_{\mathcal{P}}(\lambda,n)+\ell_{\mathcal{V}}(\lambda,n))}\cdot \mathsf{poly}(T(\lambda,n))\end{pmatrix}.$$

Proof. For each λ and $x \notin L$, apply Theorem 6.6 to the distribution $F(1^\lambda, x)$ with $\varepsilon = 2^{-\lambda}$. Since $T(\lambda, n) \geq \lambda + n$, this gives an advice mapping $\widetilde{F}$, independent of the adversary, with description complexity

$$s_{\widetilde{F}}(\lambda, n) = 2^{\ell_{\mathcal{P}}(\lambda,n)+\ell_{\mathcal{V}}(\lambda,n)}\,\mathsf{poly}(T(\lambda, n)).$$

Every target winning probability changes by at most $2^{-\lambda}$, so $(\mathcal{R}, \widetilde{F})$ satisfies the same black-box reduction guarantee.

The replacement is supported on advice strings of length at most $T(\lambda, n)$. Thus the hardness hypothesis is unchanged: the completeness argument and the hypothesis of Lemma 6.2, both consequences of security against $\mathcal{R} \circ \mathcal{P}$, bound the winning probability for each advice string separately before averaging. Applying Corollary 6.4 with the above bound on $s_{\widetilde{F}}$ gives the claimed bounds. □

6.3 Non-Uniform Reductions for Adaptive Soundness

We now consider adaptively sound SNARGs as in Definition 4.7. The reduction and its advice depend only on the security parameter and instance length. We track the advice description complexity s_F separately from the original reduction time T. This gives a laconic proof whose first message is independent of the instance, to which we can apply Theorem 5.2.

Theorem 6.8. *Fix a nondecreasing polynomial p and a security-parameter function $\lambda = \lambda(n)$ such that $n \leq p(\lambda(n))$ for every n. Suppose L has an adaptively sound SNARG $\Pi = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ whose soundness admits a non-uniform, one-query, $T(\lambda, n)$-time black-box reduction $(\mathcal{R}, F)$ to a game-based assumption $(\mathcal{C}, c)$ and advice of description complexity $s_F(\lambda, n)$. Here $\ell_{\mathcal{P}}(\lambda, n) \geq 1$ and $\ell_{\mathcal{V}}(\lambda, n)$ bound the proof and proving-key lengths.*

Assuming that $(\mathcal{C}, c)$ is secure against $\mathcal{R} \circ \mathcal{P}$ at the reduction's target-parameter map f (Definition 6.1), there is a one-round private-coin proof for L with a uniform verifier given a sampler description D_n satisfying

$$D_n(\rho) = F(1^{\lambda(n)}, 1^n; \rho), \qquad |D_n| \leq s_F(\lambda(n), n).$$

Its verifier time is $\mathsf{poly}(T(\lambda, n), s_F(\lambda, n))$, its prover and verifier message lengths are at most $\ell_{\mathcal{P}}(\lambda, n)$ and $\ell_{\mathcal{V}}(\lambda, n)$, and its error bounds satisfy

$$1 - \epsilon_c(n) - \epsilon_s(n) \geq \frac{1}{\mathsf{poly}(\lambda, n)}.$$

Completeness holds for every candidate sampler description D of length at most $s_F(\lambda, n)$, and statistical soundness holds on every $x \notin L$ of length n for the same description D_n. The first verifier message depends on x only through its length: for every allowed D and random tape ρ,

$$\mathcal{V}_1'(1^\lambda, x, D; \rho) = \mathcal{V}_1'(1^\lambda, 1^n, D; \rho).$$

Proof. Apply Lemma 6.2, with $\Gamma(\lambda, n) = \lambda$, to the adaptive soundness games of Definition 4.7, whose public index is 1^n. Thus there is an inverse-polynomial bound $\delta(\lambda, n) > 0$ such that, for all sufficiently large λ, every adaptive adversary with soundness advantage at least $1/2$ satisfies

$$\Pr\left[\mathcal{R}^{\mathcal{P}^*}(1^\lambda, 1^n, F(1^\lambda, 1^n)) \text{ wins } \mathcal{C}(1^{\lambda'})\right] \geq c + \delta(\lambda, n), \tag{10}$$

where $\lambda' = f(1^\lambda, n)$. The probability includes the independently sampled advice $F(1^\lambda, 1^n)$. We may take $\delta \leq (1-c)/2$. As in the non-adaptive proof, we assume without loss of generality that $\mathcal{R}$ never queries strings longer than $\ell_{\mathcal{V}}(\lambda, n)$.

Choose a constant $\beta \in (2c/(1+2c), 1)$ and put $\tau = 1 - c(1-\beta)$. We use the protocol from Theorem 6.3, with the following changes.

> On common input $x \in \{0,1\}^n$, the verifier has a sampler description D with $|D| \leq s_F(\lambda, n)$ and rejects prover replies longer than $\ell_{\mathcal{P}}(\lambda, n)$. Malformed descriptions and outputs longer than $T(\lambda, n)$ are interpreted as the empty string.
>
> 1. With probability β, the verifier samples $(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{G}(1^\lambda, 1^n)$. Otherwise, it samples fresh private coins ρ_F, sets $z = D(\rho_F)$, and simulates $\mathcal{R}(1^\lambda, 1^n, z)$ interacting with $\mathcal{C}(1^{\lambda'})$ using independent coins until the oracle query pk. It sends pk, keeping the branch and private state. If there is no query, it sends a fixed dummy message and records the outcome.
> 2. The prover returns $\pi \leftarrow \mathcal{P}(\mathsf{pk}, x, w)$.
> 3. In the first branch, the verifier outputs $\mathcal{V}(\mathsf{vk}, x, \pi)$. In the second branch, it supplies (x, π) as the oracle reply and resumes the simulation, accepting exactly when $\mathcal{C}$ does not output win. If there was no query, it uses the recorded outcome.

Completeness. For every $x \in L$, witness w, and z with $|z| \leq T(\lambda, n)$, the assumption adversary with x, w, z hardwired can answer the reduction's query with $(x, \mathcal{P}(\mathsf{pk}, x, w))$. This adversary is $(\mathcal{R} \circ \mathcal{P})(1^n, z, (x, w))$ (Definition 6.1), so security against $\mathcal{R} \circ \mathcal{P}$ gives a common negligible bound $\nu(\lambda)$ on its advantage at λ', which we increase to cover the original completeness error. Averaging over $z = D(\rho_F)$, the mixed verifier therefore accepts with probability at least $\tau - \nu(\lambda)$ for every allowed D.

Statistical soundness. Fix $x \notin L$ of length n and $D = D_n$, so that $D(\rho_F)$ has distribution $F(1^\lambda, 1^n)$. Given any possibly unbounded prover $\mathcal{P}^*$ for the new proof, define the adaptive adversary $\widetilde{\mathcal{P}}(\mathsf{pk}) = (x, \mathcal{P}^*(\mathsf{pk}))$. Let α_{cheat} be its adaptive soundness advantage. The mixed acceptance probability is

$$\beta\alpha_{\text{cheat}} + (1-\beta)\left(1 - \Pr\left[\mathcal{R}^{\widetilde{\mathcal{P}}}(1^\lambda, 1^n, F(1^\lambda, 1^n)) \text{ wins } \mathcal{C}(1^{\lambda'})\right]\right).$$

If $\alpha_{\text{cheat}} < 1/2$, this is at most $\beta/2 + 1 - \beta$. Otherwise, apply (10). The same calculation as in Theorem 6.3 bounds acceptance by $\tau - \xi(\lambda, n)$, where

$$\xi(\lambda, n) = \min\left\{\frac{\beta}{2} - (1-\beta)c,\ (1-\beta)\delta(\lambda, n)\right\}.$$

The sampler description is the same for every x and every prover at this length.

Choose a cutoff so that $\nu(\lambda) \leq \xi(\lambda, n)/2$ for all larger λ and all $n \leq p(\lambda)$. Set

$$\epsilon_c(n) = 1 - \tau + \frac{\xi(\lambda(n), n)}{2}, \qquad \epsilon_s(n) = \tau - \xi(\lambda(n), n).$$

These have gap $\xi/2$. On the finitely many smaller input lengths, the verifier decides membership directly. In every case, the first message is generated using only λ, n, D and the verifier's coins. Evaluating the sampler and running the mixed verifier takes time $\mathsf{poly}(T(\lambda,n), s_F(\lambda,n))$. □

Corollary 6.9. *Under the hypotheses of Theorem 6.8, writing $\lambda = \lambda(n)$, we have*

$$L \in \mathsf{SIZE}\Big[2^{\ell_{\mathcal{P}}(\lambda,n)}\,\mathsf{poly}(T(\lambda,n), s_F(\lambda,n))\Big].$$

Moreover, L can be decided deterministically in this time with advice of length $\mathsf{poly}(T(\lambda,n), s_F(\lambda,n))$ depending only on n.

Proof. Fix the sampler description D_n from Theorem 6.8. Its one-round proof has verifier time $\mathsf{poly}(T, s_F)$, an instance-independent first message, and an inverse-polynomial gap. Apply the sampling argument of Theorem 5.2 with D_n fixed, and include D_n in the resulting advice. The enumerated reply is the $\ell_{\mathcal{P}}$-bit proof π; the instance x in the reduction's oracle answer is supplied by the algorithm's input. □

Multi-query reductions. Using a different argument, based on that of Gentry and Wichs [GW11], Corollary 6.9 extends to reductions making many adaptive queries, with advice of any description complexity sampled independently of the cheating prover. We prove the following theorem in Appendix A.

Theorem 6.10. *Fix a nondecreasing polynomial p and a security-parameter function $\lambda = \lambda(n)$ such that $n \le p(\lambda(n))$ for every n. Suppose L has an adaptively sound SNARG $\Pi = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ with proof length $\ell_{\mathcal{P}}(\lambda, n)$, whose soundness admits a non-uniform, $k(\lambda,n)$-query, $T(\lambda,n)$-time black-box reduction $(\mathcal{R}, F)$ with adaptive queries to a game-based assumption $(\mathcal{C}, c)$. Assume that $(\mathcal{C}, c)$ is secure against $\mathcal{R} \circ \mathcal{P}$ at the reduction's target-parameter map f (Definition 6.1). Then, writing $\lambda = \lambda(n)$,*

$$L \in \mathsf{SIZE}\Big[2^{O(\ell_{\mathcal{P}}(\lambda,n))} \cdot \mathsf{poly}\big(T(\lambda,n)\big)\Big].$$

7 Implications of Black-Box Reductions

7.1 Implications for Witness Encryption and Obfuscation

Witness encryption.

Witness encryption gives two-message laconic arguments (see, e.g., [Wu18, Section 6.1.3]). Let $(\mathsf{Enc}, \mathsf{Dec})$ be a witness encryption scheme (Definition 4.8) for an NP relation, supporting one-bit messages and allowing negligible decryption error. On common input x, the verifier samples $\mu \leftarrow \{0,1\}$, sends $\mathsf{ct} \leftarrow \mathsf{Enc}(1^\lambda, x, \mu)$, and keeps μ private. Given a witness w, the prover returns $\pi = \mathsf{Dec}(\mathsf{ct}, w)$, and the verifier accepts if $\pi = \mu$. The prover sends exactly one bit, and the completeness error is at most the decryption error.

The soundness game is exactly the fixed-message witness-encryption game of Section 4.1.4, with baseline $c_0 = 1/2$. Indeed, on $x \notin L$, the cheating prover receives an encryption of a uniform bit μ and wins precisely when it returns μ. Thus the reduction and its query complexity are preserved. This gives soundness error at most $1/2 + \mathsf{negl}(\lambda)$, which is sufficient for Theorem 6.3.

Corollary 7.1. *Let $L \in \mathsf{NP}$ have an NP relation $\mathcal{R}$. Suppose witness encryption for $\mathcal{R}$ admits a one-query black-box security reduction to the T-hardness of a game-based assumption $(\mathcal{C}, c)$ at its target-parameter map f, and that $(\mathcal{C}, c)$ is T-hard at f (Definition 4.12), for some $T(\lambda,n) = \mathsf{poly}(\lambda, n)$. If the reduction runs*

in polynomial time and is either uniform or non-uniform with advice of description complexity $s_F(\lambda, n) = \mathsf{poly}(\lambda, n)$*, then*

$$L \in \mathsf{coAM} \begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & \mathsf{poly}(n) \\ \mathsf{prover\ comm} & = & \mathsf{poly}(n) \\ \mathsf{verifier\ comm} & = & \mathsf{poly}(n) \\ \mathsf{verifier\ randomness} & = & \mathsf{poly}(n) \end{pmatrix}.$$

Proof. As discussed above, the reduction is also a one-query soundness reduction for the one-bit argument for L. Apply Corollary 6.4 with $\Gamma = T$, $c_0 = 1/2$, $\ell_{\mathcal{P}}(\lambda, n) = 1$, $\lambda(n) = n$, and $p(\lambda) = \lambda$. Since T is polynomial, negligible decryption error satisfies the completeness condition. The verifier time is $2\,\mathsf{poly}(T(n,n), s_F(n,n)) = \mathsf{poly}(n)$, and the communication and randomness are polynomial as well. □

Corollary 7.2. *Let* $L \in \mathsf{NP}$ *have an* NP *relation* $\mathcal{R}$*. Suppose witness encryption for* $\mathcal{R}$ *admits a one-query black-box security reduction with deterministic advice, running in time* $T(\lambda, n)$*, to the* T*-hardness of a game-based assumption* $(\mathcal{C}, c)$*. Assume that, for some constant* $\theta > 0$*, the target assumption is* T*-hard when* $\lambda(n) = n^{\theta}$*.*[14] *Then*

$$L \in \mathsf{coAM} \begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & \mathsf{poly}(T(\lambda(n), n)) \\ \mathsf{prover\ comm} & = & \mathsf{poly}(T(\lambda(n), n)) \\ \mathsf{verifier\ comm} & = & \mathsf{poly}(T(\lambda(n), n)) \\ \mathsf{verifier\ randomness} & = & \mathsf{poly}(T(\lambda(n), n)) \end{pmatrix}.$$

Proof. The argument has soundness baseline $c_0 = 1/2$ and prover-message length $\ell_{\mathcal{P}}(\lambda, n) = 1$. After boosting correctness to $T^{-\omega(1)}$ (footnote to Definition 4.8), which preserves one-query reductions to T-hardness and changes T only by a $\mathsf{poly}(\log T)$ factor, apply Corollary 6.5 with $\Gamma = T$, the stated $\lambda(n)$, and $p(t) = t^{O(1/\theta)}$. The factor $2^{\ell_{\mathcal{P}}}$ is constant, giving verifier time $\mathsf{poly}(T(\lambda(n), n))$. □

Indistinguishability obfuscation.

The standard construction of witness encryption from iO [GGH+13] gives the following consequence.

As in Definition 4.11, we express iO security as an indexed family of two-message games. The public index is a pair (C_0, C_1) of equal-size, functionally equivalent circuits. The challenger chooses $b \leftarrow \{0,1\}$, sends an obfuscation of C_b, and outputs win precisely when the adversary's reply equals b; the baseline is $1/2$. Circuit equivalence is imposed through the admissible index set.

Corollary 7.3. *Suppose* iO *for circuits admits a one-query black-box security reduction to the* T*-hardness of a game-based assumption* $(\mathcal{C}, c)$ *at its target-parameter map* f*, and that* $(\mathcal{C}, c)$ *is* T*-hard at* f *(Definition 4.12), for some* $T(\lambda, n) = \mathsf{poly}(\lambda, n)$*. If the reduction runs in polynomial time and is either uniform or non-uniform with advice of polynomial description complexity, then every* $L \in \mathsf{NP}$ *satisfies the conclusion of Corollary 7.1, that is,* $\mathsf{NP} \subseteq \mathsf{coAM}$*.*

Proof. Witness encryption for NP can be constructed from iO with a one-query reduction [GGH+13]: to encrypt a bit b under x, output $\mathsf{iO}(1^{\lambda}, C_{x,b})$, where

$$C_{x,b}(w) = \begin{cases} b, & (x,w) \in \mathcal{R}, \\ \bot, & \text{otherwise}, \end{cases}$$

[14] Hardness is at the reduction's target-parameter map f, with $p(t) = t^{O(1/\theta)}$ in Definition 4.12.

padded to a size $s(n) = \mathsf{poly}(n)$. If $x \notin L$, the circuits $C_{x,0}$ and $C_{x,1}$ are functionally equivalent, and the reduction forwards the challenge obfuscation to the witness-encryption adversary. Composing this with the assumed reduction for iO gives a one-query reduction for the argument of Section 7.1, with advice $F(1^\lambda, (C_{x,0}, C_{x,1}))$, target-parameter map $(1^\lambda, n) \mapsto f(1^\lambda, s(n))$, and polynomial overhead; it is still a reduction to the T-hardness of $(\mathcal{C}, c)$. The composed honest-prover reduction has polynomial size, also on true statements, where the pair is not an admissible index, so T-hardness gives security against it. Apply Corollary 7.1 to every $L \in \mathsf{NP}$. □

For iO, the natural time bound of a reduction depends on the input length ℓ of the obfuscated circuits, which the input-length barrier compares with 2^ℓ. In the construction above, the circuits $C_{x,b}$ have input length at most $q(n)$ and size $\mathsf{poly}(n)$.

Corollary 7.4. *Suppose* iO *for circuits admits a one-query black-box security reduction with deterministic advice to the T-hardness of a game-based assumption $(\mathcal{C}, c)$, running in time $T_{\mathsf{iO}}(\lambda, \ell, s)$ on circuit pairs of input length ℓ and size s, where T_{iO} is nondecreasing in ℓ and s. Let $L \in \mathsf{NP}$ have witness bound q, let $s(n) = \mathsf{poly}(n)$ bound the size of the padded circuits $C_{x,b}$, and let $T(\lambda, n) = T_{\mathsf{iO}}(\lambda, q(n), s(n)) \cdot \mathsf{poly}(\lambda, n)$. Assuming the security required by Corollary 7.2 for the composed soundness reduction below, L satisfies the conclusion of Corollary 7.2.*

Proof. The composed soundness reduction in the proof of Corollary 7.3 makes one query, preserves deterministic advice, and adds only $\mathsf{poly}(\lambda, n)$ overhead. On a statement of length n, it invokes the iO reduction on circuits of input length at most $q(n)$ and size at most $s(n)$, so it runs in time at most $T(\lambda, n)$. It is still a reduction to the T-hardness of $(\mathcal{C}, c)$, since the composition with $x \mapsto (C_{x,0}, C_{x,1})$ preserves the advantage. Apply Corollary 7.2; correctness is boosted in the same way (footnote to Definition 4.11). □

7.2 Implications for SNARGs

7.2.1 Non-Adaptive Soundness

A SNARG gives a two-message laconic argument by taking the proving key as the verifier's first message: the verifier samples $(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{G}(1^\lambda, 1^{|x|})$, sends pk to the prover, and checks the returned proof π using $\mathcal{V}(\mathsf{vk}, x, \pi)$. This preserves the proof length, completeness, and non-adaptive soundness, as well as the soundness reduction and its query complexity.

Corollary 7.5. *Suppose $L \in \mathsf{NP}$ has a SNARG whose non-adaptive soundness admits a polynomial-time, one-query black-box reduction to a game-based assumption that is T-hard at the reduction's target-parameter map f (Definition 4.12) for some $T(\lambda, n) = \mathsf{poly}(\lambda, n)$. If the reduction is either uniform or non-uniform with advice of description complexity $s_F(\lambda, n) = \mathsf{poly}(\lambda, n)$, then, for every constant $\epsilon > 0$,*

$$L \in \mathsf{coAM} \begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & 2^{n^\epsilon} \\ \mathsf{prover\ comm} & = & \mathsf{poly}(n) \\ \mathsf{verifier\ comm} & = & \mathsf{poly}(n) \\ \mathsf{verifier\ randomness} & = & \mathsf{poly}(n) \end{pmatrix}.$$

Proof. Fix an NP relation for L with witness bound $q(n) = \mathsf{poly}(n)$, and write the SNARG proof-length bound as $\ell_{\mathcal{P}}(\lambda, n) = O(\lambda^\alpha)(n + q(n))^{o(1)}$ for a constant $\alpha > 0$. Viewing the proving key as the verifier's first message, apply Corollary 6.5 or Corollary 6.4 with $T(\lambda, n) = \mathsf{poly}(\lambda, n)$ and, in the non-uniform case, $s_F(\lambda, n) = \mathsf{poly}(\lambda, n)$.

For a given $\epsilon > 0$, set $\lambda(n) = n^\theta$ for a sufficiently small constant $\theta = \Theta(\epsilon/\alpha)$, and take $p(t) = t^{O(1/\theta)}$. Then $n \le p(\lambda)$, and every witness has length $q(n) = \lambda^{O(1)}$. T-hardness at f therefore implies security against $\mathcal{R} \circ \mathcal{P}$ (Definition 6.1), which is all the two corollaries require. Since q is polynomial, the prover-message bound becomes $\ell_{\mathcal{P}}(\lambda(n), n) = n^{\alpha\theta+o(1)}$. The resulting verifier runs in time $2^{n^{\alpha\theta+o(1)}}\,\mathsf{poly}(n) \le 2^{n^\epsilon}$ for all sufficiently large n, and the communication and randomness remain polynomial in n. □

Remark 7.6. *Chen and Jin [CJ25] also consider reductions that make $k = k(\lambda)$ non-adaptive queries, where k is independent of the statement length. There, the reduction's oracle access is viewed as one ordered batched query $(\mathsf{pk}_1, \ldots, \mathsf{pk}_t)$ with $t \le k$, answered jointly by the cheating prover, while soundness is measured on batches of size one. Such a reduction is a one-query reduction for the following laconic argument $\Pi^{\le k}$: the verifier sends a batch (pk) of size one, with $(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x)$, and applies $\mathcal{V}_2(\mathsf{vk}, x, \cdot)$ to the first proof of the reply. A cheating prover for $\Pi^{\le k}$ is any function from batches of size at most k to lists of proofs, which is exactly the class of adversaries in [CJ25], and its soundness advantage is its non-adaptive soundness advantage against Π. The argument $\Pi^{\le k}$ has prover-message length $k\ell_{\mathcal{P}}$ and verifier-message length $k\ell_{\mathcal{V}}$. Since k depends only on λ, we have $k\ell_{\mathcal{P}} = \mathsf{poly}(\lambda) \cdot (n + q(n))^{o(1)}$, and the proof of Corollary 7.5 applies unchanged. In particular, our results also cover this setting, with no restriction on the CRS length.*

We next allow larger reduction times. The reduction still achieves nonnegligible advantage, as in Definition 4.13. We require the target assumption to be T-hard at the reduction's target-parameter map, which covers the simulations used in the proof, including their polynomial overhead.

Corollary 7.7. *Let $\alpha, \theta > 0$ be constants, and set $\lambda(n) = n^\theta$. Suppose $L \in \mathsf{NP}$ has a SNARG with proof length $\ell_{\mathcal{P}}(\lambda, n) = O(\lambda^\alpha)(n + q(n))^{o(1)}$, whose non-adaptive soundness admits a one-query black-box reduction with deterministic advice, running in time $T(\lambda, n)$, to a game-based assumption $(\mathcal{C}, c)$. Assume that $(\mathcal{C}, c)$ is T-hard at the reduction's target-parameter map f (Definition 4.12), for statement lengths $n \le \lambda^{O(1/\theta)}$. Then*

$$L \in \mathsf{coAM}\left(\begin{array}{lcl} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & 2^{n^{\alpha\theta+o(1)}} \cdot \mathsf{poly}(T(\lambda(n), n)) \\ \mathsf{prover\ comm} & = & \mathsf{poly}(T(\lambda(n), n)) \\ \mathsf{verifier\ comm} & = & \mathsf{poly}(T(\lambda(n), n)) \\ \mathsf{verifier\ randomness} & = & \mathsf{poly}(T(\lambda(n), n)) \end{array}\right).$$

In particular, for every constant $\epsilon > 0$, choosing a sufficiently small $\theta = \Theta(\epsilon/\alpha)$ gives verifier time $2^{n^\epsilon} \cdot \mathsf{poly}(T(\lambda(n), n))$.

Proof. Apply Corollary 6.5, which covers uniform reductions as the special case of empty advice, with $\lambda(n) = n^\theta$ and $p(t) = t^{O(1/\theta)}$. Then $n \le p(\lambda(n))$, and the assumed security is exactly the hardness hypothesis of that corollary. Since q is polynomial, the prover-message bound becomes $\ell_{\mathcal{P}}(\lambda(n), n) = O(n^{\alpha\theta}) \cdot n^{o(1)} = n^{\alpha\theta+o(1)}$, and the claimed bounds follow. If $\alpha\theta < \epsilon$, then $n^{\alpha\theta+o(1)} \le n^\epsilon$ for all sufficiently large n. □

For example, if $T(\lambda, n) = 2^{O(n^\beta)}\,\mathsf{poly}(\lambda)$ for a constant $\beta > 0$, then the corollary gives verifier time $2^{n^\epsilon} \cdot 2^{O(n^\beta)}$ and communication $2^{O(n^\beta)}$. The required security is against adversaries of size $2^{O(n^\beta)} = 2^{O(\lambda^{\beta/\theta})}$ at the target parameter $\lambda' = f(1^\lambda, n)$, where $\lambda = n^\theta$. For $f(1^\lambda, n) = \lambda$, this is subexponential security with exponent β/θ. For polynomial T, Corollary 7.5 is the special case with polynomial communication.

7.2.2 Adaptive Soundness

Corollary 7.8. *Suppose $L \in \mathsf{NP}$ has an adaptively sound SNARG with a $T(\lambda, n)$-time, multi-query black-box soundness reduction with adaptive queries and possibly randomized advice to a game-based assumption $(\mathcal{C}, c)$, with target-parameter map f. For every constant $\epsilon > 0$, there is a constant $\theta > 0$ such that, setting $\lambda(n) = n^\theta$, T-hardness of $(\mathcal{C}, c)$ at f[15] implies*

$$L \in \mathsf{SIZE}\Big[2^{n^\epsilon} \cdot \mathsf{poly}\big(T(\lambda(n), n)\big)\Big].$$

If the reduction makes one query and its advice has description complexity $\mathsf{poly}(T(\lambda, n))$, then L also has a deterministic algorithm with the same time bound and advice of length $\mathsf{poly}(T(\lambda(n), n))$.

Proof. Write $\ell_{\mathcal{P}}(\lambda, n) = O(\lambda^\alpha)(n + q(n))^{o(1)}$ for a constant $\alpha > 0$, and choose a sufficiently small constant $\theta = \Theta(\epsilon/\alpha)$. Since $\ell_{\mathcal{P}}(\lambda(n), n) \le n^{\alpha\theta + o(1)}$, Theorem 6.10 gives circuits of size

$$2^{O(\ell_{\mathcal{P}}(\lambda(n), n))}\, \mathsf{poly}(T(\lambda(n), n)) \le 2^{n^\epsilon}\, \mathsf{poly}(T(\lambda(n), n))$$

for all sufficiently large n. For one-query reductions, apply Corollary 6.9 with $s_F(\lambda, n) = \mathsf{poly}(T(\lambda, n))$ to obtain the claimed time and advice bounds. □

7.3 Implication for Adversary-Dependent Advice

We now allow the advice of a one-query reduction for a SNARG to depend on the cheating prover. For a short proving key, this recovers the one-query lower bound of [CJ25] from a worst-case hypothesis.

Definition 7.9 (Black-box reduction with adversary-dependent advice)**.** *A black-box reduction for the non-adaptive soundness of a SNARG* with adversary-dependent advice *is defined as in Definition 4.13, except that the advice is sampled as $z \leftarrow F(1^\lambda, x, \mathcal{P}^*)$, where $\mathcal{P}^*$ is the possibly unbounded cheating prover. We still require $|z| \le T(\lambda, |x|)$.*

Corollary 7.10. *Suppose $L \in \mathsf{NP}$ has a SNARG with proof length $\ell_{\mathcal{P}}(\lambda, n)$ and proving-key length $\ell_{\mathcal{V}}(\lambda, n)$, whose non-adaptive soundness admits a T-time, one-query black-box reduction with adversary-dependent advice to a game-based assumption that is T-hard at the reduction's target-parameter map (Definition 4.12). Let $\lambda = \lambda(n)$ satisfy $n \le p(\lambda(n))$ for a polynomial p. Then*

$$L \in \mathsf{coAM}\begin{pmatrix} \mathsf{rounds} & = & 1 \\ \mathsf{verifier\ time} & = & 2^{O(\ell_{\mathcal{P}} + \ell_{\mathcal{V}})} \cdot \mathsf{poly}(T) \\ \mathsf{prover\ comm} & = & 2^{O(\ell_{\mathcal{P}} + \ell_{\mathcal{V}})} \cdot \mathsf{poly}(T) \\ \mathsf{verifier\ comm} & = & 2^{O(\ell_{\mathcal{P}} + \ell_{\mathcal{V}})} \cdot \mathsf{poly}(T) \\ \mathsf{verifier\ randomness} & = & 2^{O(\ell_{\mathcal{P}} + \ell_{\mathcal{V}})} \cdot \mathsf{poly}(T) \end{pmatrix},$$

where $\ell_{\mathcal{P}}, \ell_{\mathcal{V}}, T$ are evaluated at $(\lambda(n), n)$. In particular, if T is polynomial and $\ell_{\mathcal{P}}(\lambda, n) + \ell_{\mathcal{V}}(\lambda, n) = O(\lambda^\alpha) \cdot n^{o(1)}$ for a constant $\alpha > 0$, then $L \in \mathsf{coAM}[2^{n^\epsilon}]$ for every constant $\epsilon > 0$.

Proof. Viewing the proving key as the verifier's first message, the SNARG gives a two-message laconic argument with message lengths $\ell_{\mathcal{V}}$ and $\ell_{\mathcal{P}}$, and the reduction is a one-query reduction for its non-adaptive soundness (Section 7.2.1). We show that the adversary-dependent advice can be replaced by advice that is independent of the cheating prover and has description complexity

[15] We use Definition 4.12 with $p(t) = t^{O(1/\theta)}$.

$2^{\ell_{\mathcal{P}}+\ell_{\mathcal{V}}}\,\mathsf{poly}(T)$. Corollary 6.4 then gives the claimed bounds, and the particular case follows by choosing $\lambda(n)$ as in the proof of Corollary 7.5.

Completeness in Theorem 6.3 and the hypothesis of Lemma 6.2 bound the winning probability of $\mathcal{R}^{\mathcal{P}}$ for each advice string separately, so they are unaffected by how the advice is chosen. It remains to obtain the conclusion of Lemma 6.2, with $\eta = (1-c_0)/2$, for a single advice distribution that works against all cheating provers.

A common bound for the best advice string. As in the proof of Lemma 6.2, there is a polynomial q such that, for all sufficiently large λ, every $x \notin L$ with $|x| \le p(\lambda)$, and every cheating prover $\mathcal{P}^*$ with soundness advantage at least $(1-c_0)/2$, some advice string z with $|z| \le T(\lambda, |x|)$ satisfies

$$\Pr\big[\mathcal{R}^{\mathcal{P}^*}(1^\lambda, x, z) \text{ wins } \mathcal{C}(1^{f(1^\lambda,|x|)})\big] \ge c + \frac{1}{q(\lambda,|x|)}. \tag{11}$$

Indeed, if no such q existed, the diagonalization in that proof would select provers for which (11) fails for *every* advice string, so the reduction has negligible advantage against them for any choice of advice, including adversary-dependent advice.

Removing the dependence on the cheating prover. Fix λ and $x \notin L$ with $n = |x| \le p(\lambda)$, and let Z be the set of advice strings of length at most $T(\lambda, n)$. Represent a cheating prover $\mathcal{P}^*$ by the vector of probabilities $\Pr[\mathcal{P}^*(\mathsf{pk}) = \pi]$ for $\mathsf{pk} \in \{0,1\}^{\ell_{\mathcal{V}}}$ and $\pi \in \{0,1\}^{\ell_{\mathcal{P}}}$. Since $\mathcal{R}$ makes one query, both its winning probability $W(z, \mathcal{P}^*)$ with advice z and the acceptance probability of $\mathcal{P}^*$ are linear in this vector. Thus the set $\mathcal{P}$ of cheating provers with soundness advantage at least $(1-c_0)/2$ is compact and convex, and by the minimax theorem [Sio58] and (11),

$$\max_{\mathcal{D}} \min_{\mathcal{P}^*\in\mathcal{P}} \mathbb{E}_{z\leftarrow\mathcal{D}}\big[W(z,\mathcal{P}^*)\big] = \min_{\mathcal{P}^*\in\mathcal{P}} \max_{z\in Z} W(z,\mathcal{P}^*) \ge c + \frac{1}{q(\lambda,n)},$$

where $\mathcal{D}$ ranges over distributions on Z. A maximizing distribution $\mathcal{D}_{\lambda,x}$ depends only on λ and x.

Bounding the description complexity. Applying Theorem 6.6 to $\mathcal{D}_{\lambda,x}$ with $\varepsilon = 1/(2q(\lambda,n))$ gives a distribution $\widetilde{\mathcal{D}}_{\lambda,x}$ on Z, independent of the cheating prover, with a sampler of description complexity $2^{\ell_{\mathcal{P}}+\ell_{\mathcal{V}}}\,\mathsf{poly}(T)$, against which every $\mathcal{P}^* \in \mathcal{P}$ wins with probability at least $c + 1/(2q(\lambda,n))$. This is the required conclusion of Lemma 6.2. □

8 From Multi-Query Reductions to Multiprover Interactive Proofs

We first introduce a stronger notion of correctness for MIP.

Definition 8.1 (Strong local correctness)**.** *Let $\mathcal{R}$ be an* NP *relation with language L. A multiprover protocol with specified honest provers $\mathcal{P}_1, \dots, \mathcal{P}_k$, each of size* $\mathsf{poly}(n)$ *on instances of length n, has* strong local correctness *with error ϵ_c if, for every $x \in L$ and every tuple $(w_1, \dots, w_k)$ satisfying $(x, w_i) \in \mathcal{R}$ for all $i \in [k]$,*

$$\Pr\big[\langle \mathcal{P}_1(w_1), \dots, \mathcal{P}_k(w_k), \mathcal{V}\rangle(x) = 1\big] \ge 1 - \epsilon_c(|x|).$$

Here $\mathcal{P}_i(w_i)$ denotes prover i given witness w_i. Its replies depend only on x, w_i, its local transcript, and independent private coins. The honest algorithms are fixed before the witnesses are chosen, with no joint preprocessing or shared randomness. The witnesses may differ. Soundness is as in Definition 4.32.

8.1 From One Stateless Prover to Independent Provers

We convert a reduction that queries one stateless prover several times into a reduction that queries each of several independent provers at most once.

Lemma 8.2. *Let $\mathcal{D}$ be a two-message game, let $(\mathcal{C}, c)$ be a target game, and let $\eta, \delta : \mathbb{N} \to (0, 1]$. At security parameter λ, suppose a T-time reduction $\mathcal{R}$ makes at most $k \geq 1$ adaptive queries to a stateless randomized oracle prover and achieves advantage at least $\delta(\lambda)$ against $\mathcal{C}$ whenever the prover wins $\mathcal{D}$ with probability at least $\eta(\lambda)$. Each oracle call uses fresh independent coins. Public inputs and advice are fixed, and $\mathcal{C}$ uses the target parameter chosen by $\mathcal{R}$.*

Then there is a reduction $\mathcal{R}'$ with access to $k' = O(k^2/\delta(\lambda))$ independent oracle provers, each queried at most once, running in time

$$T' = O(T) + \mathsf{poly}\big(k, \lambda, \log(1/\delta(\lambda))\big),$$

such that

$$\frac{1}{k'} \sum_{i=1}^{k'} \Pr[\mathcal{P}_i^* \textit{ wins } \mathcal{D}] \geq \eta(\lambda) \quad \Longrightarrow \quad \Pr[\mathcal{R}'^{\mathcal{P}_1^*, \ldots, \mathcal{P}_{k'}^*} \textit{ wins } \mathcal{C}] \geq c + \delta(\lambda)/2.$$

Each prover sees only its query, public inputs, advice, and identity, and uses independent private coins. The reduction preserves the advice and target parameter. The transformation preserves non-adaptivity of the queries.

Proof. Fix λ, the public input, and the advice, and write $\eta = \eta(\lambda)$ and $\delta = \delta(\lambda)$. Take $k' \geq 2k$ large enough that

$$\frac{k(k-1)}{2k'} \leq \frac{\delta}{4}, \qquad k' = O(k^2/\delta).$$

First consider the simulation that, on each oracle query of $\mathcal{R}$, chooses a uniformly random unused prover, sends it the query, and returns its reply to $\mathcal{R}$. Fix a tuple $(\mathcal{P}_1^*, \ldots, \mathcal{P}_{k'}^*)$, and define a single response oracle $\mathcal{P}^*$ that, on input u, samples $I \leftarrow [k']$ uniformly and returns $\mathcal{P}_I^*(u)$, with fresh coins on each call. Since $\mathcal{D}$ sends one challenge and receives one response,

$$\Pr[\mathcal{P}^* \text{ wins } \mathcal{D}] = \frac{1}{k'} \sum_{i=1}^{k'} \Pr[\mathcal{P}_i^* \text{ wins } \mathcal{D}].$$

Thus, if the average winning probability is at least η, then $\mathcal{R}^{\mathcal{P}^*}$ wins $\mathcal{C}$ with probability at least $c+\delta$.

Couple the two executions using the same reduction and challenger coins. Until a prover identity repeats in the single-oracle execution, use the same identities and answer coins in both executions. If the next identity repeats, the simulation instead chooses a uniformly random unused identity. At query j, every unused identity is selected with probability

$$\frac{1}{k'} + \frac{j-1}{k'} \cdot \frac{1}{k'-j+1} = \frac{1}{k'-j+1},$$

so the simulation has the required distribution. The transcripts agree unless a collision occurs among at most k independent uniform identities. By a union bound, this has probability at most $k(k-1)/(2k') \leq \delta/4$. Pre-sampling all k identities gives the same bound when the queries or stopping time are adaptive.

To run in bounded time, $\mathcal{R}'$ samples unused identities by rejection sampling, truncated after $O(\lambda + \log(k/\delta))$ attempts per query. Since $k' \geq 2k$, truncation fails with probability at most $2^{-\lambda}\delta/4$, so the total loss is at most $\delta/2$. Together with simulating $\mathcal{R}$, this gives running time $T' = O(T) + \mathsf{poly}(k, \log(1/\delta))$. The simulation preserves the advice and target parameter, and non-adaptive queries remain non-adaptive. □

8.2 Construction of the MIP

Theorem 8.3. *Fix a nondecreasing polynomial p and a security-parameter function $\lambda = \lambda(n)$ such that $n \le p(\lambda(n))$ for every n. Suppose L has a two-message laconic argument $\Pi = (\mathcal{V} = (\mathcal{V}_1, \mathcal{V}_2), \mathcal{P})$ with soundness baseline c_0 for an NP relation $\mathcal{R}$, with verifier-message length $\ell_{\mathcal{V}}(\lambda, n)$ and prover-message length $\ell_{\mathcal{P}}(\lambda, n)$. Suppose its non-adaptive soundness admits a non-uniform, $k(\lambda, n)$-query, $T(\lambda, n)$-time black-box reduction $(\mathcal{R}, F)$ to the Γ-hardness of a game-based assumption $(\mathcal{C}, c)$ at its target-parameter map f (Definition 4.13), for a bound $\lambda \le \Gamma(\lambda, n) \le T(\lambda, n)$, where $1 \le k(\lambda, n) \le T(\lambda, n)$, with deterministic advice $F(1^\lambda, x)$.[16] Assume that the completeness error κ of Π satisfies, for every constant $K > 0$, $\kappa(\lambda, n) \le \Gamma(\lambda, n)^{-K}$ for all sufficiently large λ and all $n \le p(\lambda)$.*

Define the promise problem

$$L^F_{\mathrm{yes}} = \big\{(1^\lambda, x, z) : x \in L,\ |z| \le T(\lambda, n)\big\},$$
$$L^F_{\mathrm{no}} = \big\{(1^\lambda, x, F(1^\lambda, x)) : x \notin L\big\},$$

where $n = |x|$ and $\lambda = \lambda(n)$. Assume $(\mathcal{C}, c)$ is secure against $\mathcal{R} \circ \mathcal{P}$ at the reduction's target-parameter map f (Definition 6.1); this holds whenever $(\mathcal{C}, c)$ is T-hard at f. Then there exist a prover bound $k' = k'(\lambda, n)$ and completeness and soundness error bounds ϵ_c, ϵ_s satisfying

$$k'(\lambda, n) \le k(\lambda, n)^2\, \mathsf{poly}(\Gamma(\lambda, n)), \qquad 1 - \epsilon_c(n) - \epsilon_s(n) \ge \frac{1}{\mathsf{poly}(\Gamma(\lambda, n))},$$

such that

$$(L^F_{\mathrm{yes}}, L^F_{\mathrm{no}}) \in \mathsf{MIP}\begin{pmatrix} \mathsf{provers} & = & k'(\lambda, n) \\ \mathsf{verifier\ time} & = & \mathsf{poly}(T(\lambda, n)) \\ \mathsf{ind\ prover\ comm} & = & \ell_{\mathcal{P}}(\lambda, n) \\ \mathsf{verifier\ comm} & = & k(\lambda, n)\ell_{\mathcal{V}}(\lambda, n) \\ \mathsf{completeness} & \ge & 1 - \epsilon_c \\ \mathsf{soundness} & \le & \epsilon_s \end{pmatrix}.$$

The verifier is uniform and contacts at most $k(\lambda, n)$ distinct provers. The protocol has strong local correctness: for every $x \in L$, every allowed z, and every tuple $(w_1, \ldots, w_{k'})$ with $(x, w_i) \in \mathcal{R}$, the specified honest provers are accepted with probability at least $1 - \epsilon_c(n)$. Each honest prover runs in time $\mathsf{poly}(\lambda, n)$ given its own witness. Resource bounds are expressed in terms of $n = |x|$, while the common input is $(1^\lambda, x, z)$. If the reduction's queries are non-adaptive, the resulting protocol is non-adaptive as well.

Proof. We follow the mixed-verifier construction of Theorem 6.3, first applying Lemma 8.2 so that no prover receives more than one query. Throughout, suppress the arguments (λ, n) of the resource bounds. As in that theorem, we assume without loss of generality that $\mathcal{R}$ never queries strings longer than $\ell_{\mathcal{V}}$ and that all oracle answers have at most $\ell_{\mathcal{P}}$ bits.

For a randomized response oracle $\mathcal{P}^*$, write

$$\alpha_x(\mathcal{P}^*) = \Pr\left[\ \mathcal{V}_2(\mathsf{vk}, x, \pi) = 1 \ \ :\ \ \begin{array}{c}(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x) \\ \pi \leftarrow \mathcal{P}^*(\mathsf{pk})\end{array}\ \right].$$

Apply Lemma 6.2 with $\eta = (1 - c_0)/2$; its hypothesis, Γ-security against $\mathcal{R} \circ \mathcal{P}$, follows from security against $\mathcal{R} \circ \mathcal{P}$, since $\Gamma \le T$. Let $\delta = \delta(\lambda, n) = \Gamma(\lambda, n)^{-C}$, for a constant C, be the resulting advantage bound, which applies to every false statement x of length $n \le p(\lambda)$ and every randomized response

[16] Randomized advice may be supplied by a fixed circuit description of its sampler, as in Definition 4.15. The bound T then includes the time to evaluate this circuit.

oracle, for all sufficiently large λ. Applying Lemma 8.2 to the soundness game of Π, with winning-probability threshold $(1 + c_0)/2$, gives a reduction $\mathcal{R}'$ with $k' = O(k^2/\delta)$ provers that contacts at most k of them, each at most once, and satisfies

$$\frac{1}{k'}\sum_{i=1}^{k'} \alpha_x(\mathcal{P}_i^*) \geq \frac{1+c_0}{2} \quad \Longrightarrow \quad \Pr\big[\mathcal{R}'^{\mathcal{P}_1^*,\ldots,\mathcal{P}_{k'}^*}(1^\lambda, x, F(1^\lambda, x)) \text{ wins } \mathcal{C}(1^{\lambda'})\big] \geq c + \delta/2, \tag{12}$$

where $\lambda' = f(1^\lambda, n)$. The common advantage bound is chosen before k'; it does not depend on the prover tuple.

Choose a constant $\beta \in (0, 1)$ satisfying $\beta > 2c/(1 - c_0 + 2c)$, and set $\tau = 1 - (1 - \beta)c$. The protocol is as follows.

On common input $(1^\lambda, x, z)$, the verifier $\mathcal{V}'$ rejects if $|z| > T$. Otherwise, it samples a private bit b with $\Pr[b = 0] = \beta$.

- If $b = 0$, choose $I \leftarrow [k']$ uniformly and sample $(\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x)$. Send pk to prover I, receive its reply π, and output $\mathcal{V}_2(\mathsf{vk}, x, \pi)$.
- If $b = 1$, simulate $\mathcal{R}'(1^\lambda, x, z)$ interacting with $\mathcal{C}(1^{\lambda'})$, where $\lambda' = f(1^\lambda, n)$. Whenever $\mathcal{R}'$ queries prover i on pk, send pk to that prover and return its reply to $\mathcal{R}'$. Reject if $\mathcal{C}$ outputs win in this simulation; otherwise, accept.

Each reply is restricted to at most $\ell_{\mathcal{P}}$ bits, with the same normalization as above. Honest prover $\mathcal{P}_i'$ with witness w_i returns $\mathcal{P}(\mathsf{pk}, x, w_i)$, using independent private coins and ignoring z.

Each prover is contacted at most once and receives only its question, with no additional branch, query-order, or timing information.

Strong local correctness. Fix $x \in L$, any advice z of length at most T, and any tuple of witnesses $(w_1, \ldots, w_{k'})$ with $(x, w_i) \in \mathcal{R}$. When $b = 0$, completeness of Π gives acceptance probability at least $1 - \kappa(\lambda, n)$ for every selected prover, where $\kappa \leq \Gamma^{-(C+1)}$ for all sufficiently large λ.

When $b = 1$, the verifier simulates $\mathcal{R}'(1^\lambda, x, z)$ with its queries answered by the honest provers. Condition on the identities $I_1, \ldots, I_k$ that $\mathcal{R}'$ selects for its queries, and on the event that no query exhausts its sampling attempts, which fails with probability at most $k \cdot 2^{-\lambda}\delta/(4k) = 2^{-\lambda}\delta/4$. On this event, the simulation is distributed exactly as $(\mathcal{R} \circ \mathcal{P})(x, z, w_{I_1}, \ldots, w_{I_k})$ (Definition 6.1). Thus security against $\mathcal{R} \circ \mathcal{P}$ (applied to the worst allowed inputs at each λ, using $\Gamma \leq T$) bounds the winning probability against $\mathcal{C}(1^{\lambda'})$ by $c + \Gamma^{-(C+1)} + 2^{-\lambda}\delta/4$ for all sufficiently large λ, uniformly over all allowed $x, z, w_1, \ldots, w_{k'}$. Writing $\nu = 2\Gamma^{-(C+1)} + 2^{-\lambda}\delta/4$, the verifier's acceptance probability is at least

$$\beta\big(1 - \Gamma^{-(C+1)}\big) + (1 - \beta)\big(1 - c - \Gamma^{-(C+1)} - 2^{-\lambda}\delta/4\big) \geq \tau - \nu.$$

The honest algorithms are fixed, and each uses only its own witness and independent private coins. Thus this proves the required local correctness even when the witnesses differ.

Statistical soundness. Fix a sufficiently large λ, a false statement x of length $n \leq p(\lambda)$, and the advice $z = F(1^\lambda, x)$. Fix all shared and private coins of arbitrary cheating provers. Since each prover receives at most one question and no scheduling information, the resulting prover $\mathcal{P}_i^*$ is a deterministic function of its question, with $(1^\lambda, x, F(1^\lambda, x))$ and its identity fixed.

Let α be the acceptance probability when $b = 0$. Conditioning on the uniformly selected prover

gives

$$\alpha = \Pr\left[\ \mathcal{V}_2(\mathsf{vk}, x, \pi) = 1 \ \ :\ \ \begin{array}{c} I \leftarrow [k'],\ (\mathsf{pk}, \mathsf{vk}) \leftarrow \mathcal{V}_1(1^\lambda, x) \\ \pi = \mathcal{P}_I^*(\mathsf{pk}) \end{array}\ \right] = \frac{1}{k'} \sum_{i=1}^{k'} \alpha_x(\mathcal{P}_i^*).$$

Each summand is prover i's success probability on an independently generated first message of the original argument. The verifier's total acceptance probability is

$$\beta\alpha + (1-\beta)\left(1 - \Pr\left[\mathcal{R}'^{\mathcal{P}_1^*,\ldots,\mathcal{P}_{k'}^*}(1^\lambda, x, F(1^\lambda, x)) \text{ wins } \mathcal{C}(1^{\lambda'})\right]\right). \tag{13}$$

If $\alpha < (1+c_0)/2$, bounding acceptance in the $b = 1$ branch by 1 gives

$$\Pr[\mathcal{V}' \text{ accepts}] \le \frac{\beta(1+c_0)}{2} + (1-\beta) = \tau - \left(\frac{\beta(1-c_0)}{2} - (1-\beta)c\right).$$

Otherwise, $\alpha \ge (1+c_0)/2$, so (12) applies to the fixed deterministic tuple. Substituting into (13) and using $\alpha \le 1$, we obtain

$$\Pr[\mathcal{V}' \text{ accepts}] \le \beta + (1-\beta)(1 - c - \delta/2) = \tau - (1-\beta)\delta/2.$$

Consequently, acceptance is at most $\tau - \xi$, where

$$\xi = \min\left\{\frac{\beta(1-c_0)}{2} - (1-\beta)c,\ (1-\beta)\delta/2\right\} = \Omega(\delta).$$

This bound holds for every deterministic tuple. Averaging over the shared and private prover coins preserves it, since these coins are independent of the verifier's coins.

Parameters. Increase the cutoff λ_0 so that $\nu \le \xi(\lambda, n)/2$ for all $\lambda \ge \lambda_0$ and $n \le p(\lambda)$; this is possible since $\xi = \Omega(\delta)$, $\delta = \Gamma^{-C}$, and $\Gamma \ge \lambda$. Set

$$\epsilon_c(n) = 1 - \tau + \xi(\lambda, n)/2, \qquad \epsilon_s(n) = \tau - \xi(\lambda, n).$$

The gap is at least $\xi/2 = 1/\operatorname{poly}(\Gamma(\lambda, n))$. On the finitely many lengths with $\lambda(n) < \lambda_0$, the verifier decides membership directly.

The verifier receives z explicitly and never computes $F(1^\lambda, x)$. It contacts at most k provers, each once, and each prover sends at most $\ell_{\mathcal{P}}$ bits. The total verifier communication is at most $k\ell_{\mathcal{V}}$. Its running time is $\operatorname{poly}(T)$, including the bounded-time simulation of $\mathcal{R}'$ and $\mathcal{C}$. Each honest prover runs the original polynomial-time algorithm on its own witness. Finally, if $\mathcal{R}$ chooses its queries non-adaptively, so does $\mathcal{R}'$, and both branches of the protocol are non-adaptive. □

For uniform reductions, the advice is empty, and Theorem 8.3 gives an MIP for L itself. Choosing $\lambda(n)$ as a small polynomial in n makes each reply short.

Corollary 8.4. *Suppose L has a two-message laconic argument with soundness baseline c_0 for an* NP *relation $\mathcal{R}$ with prover-message length $\ell_{\mathcal{P}}(\lambda, n) \le \lambda^d \cdot n^{o(1)}$ for a constant $d \ge 1$. Suppose its non-adaptive soundness admits a uniform, $k(\lambda, n)$-query, $T(\lambda, n)$-time black-box reduction to a game-based assumption $(\mathcal{C}, c)$. Assume that $(\mathcal{C}, c)$ is secure against $\mathcal{R} \circ \mathcal{P}$ at the reduction's target-parameter map (Definition 6.1) for every polynomial bound on the statement length; this holds if $(\mathcal{C}, c)$ is T-hard at that map. Then, for every constant $\epsilon' > 0$, setting $\lambda(n) = \lceil n^{\epsilon'/(2d)} \rceil$ and $\widetilde{T}(n) = T(\lambda(n), n)$, L has an MIP with strong local correctness and $\operatorname{poly}(\widetilde{T}(n))$ provers, in which the verifier runs in time $\operatorname{poly}(\widetilde{T}(n))$ and contacts at most $k(\lambda(n), n) \le \widetilde{T}(n)$ provers, each at most once, each reply has at most $n^{\epsilon'}$ bits, and $1 - \epsilon_c(n) - \epsilon_s(n) \ge 1/\operatorname{poly}(n)$. Each honest prover runs in time $\operatorname{poly}(n)$ given its own witness. If the reduction's queries are non-adaptive, the resulting protocol is non-adaptive as well.*

Proof. WLOG $\epsilon' \le 1$. Apply Theorem 8.3 with $\lambda(n) = \lceil n^{\epsilon'/(2d)} \rceil$ and $p(\lambda) = \lambda^{\lceil 2d/\epsilon' \rceil}$, so that $n \le p(\lambda(n))$. Since the reduction is uniform, F outputs the empty string, and the verifier computes $1^{\lambda(n)}$ from $|x|$. Thus, the protocol is an MIP for L. Each reply has at most $\ell_{\mathcal{P}}(\lambda(n), n) \le \lambda(n)^d \cdot n^{o(1)} \le n^{\epsilon'/2+o(1)} \le n^{\epsilon'}$ bits for sufficiently large n. Since $\lambda(n) \le n+1$, every quantity bounded by a polynomial in (λ, n) is bounded by a polynomial in n, which gives the bounds on the gap and the honest provers. The verifier runs in time $\mathsf{poly}(\widetilde{T}(n))$, and the number of provers is $k^2\,\mathsf{poly}(\lambda, n) \le \mathsf{poly}(\widetilde{T}(n))$, since $k \le T$ and $T(\lambda, n) \ge \lambda + n$. The hardness hypothesis of Theorem 8.3 is the assumed security against $\mathcal{R} \circ \mathcal{P}$, with $p(\lambda) = \lambda^{2d/\epsilon'}$.

□

Corollary 8.5. *Under the hypotheses of Theorem 8.3, suppose the reduction is uniform, $\lambda(n) = n$, and $\ell_{\mathcal{P}}(\lambda, n) = 1$. Then there are a prover bound $k'(n) \le k(n,n)^2\,\mathsf{poly}(n)$ and error bounds ϵ_c, ϵ_s with*

$$1 - \epsilon_c(n) - \epsilon_s(n) \ge \frac{1}{\mathsf{poly}(n)}$$

such that

$$L \in \mathsf{MIP}\begin{pmatrix} \mathsf{provers} & = & k'(n) \\ \mathsf{verifier\ time} & = & \mathsf{poly}(T(n,n)) \\ \mathsf{ind\ prover\ comm} & = & 1 \\ \mathsf{verifier\ comm} & = & k(n,n)\ell_{\mathcal{V}}(n,n) \\ \mathsf{completeness} & \ge & 1-\epsilon_c \\ \mathsf{soundness} & \le & \epsilon_s \end{pmatrix}.$$

The protocol has strong local correctness, and each honest prover runs in time $\mathsf{poly}(n)$ given its own witness. The verifier contacts at most $k(n,n)$ distinct provers, each at most once. If the reduction's queries are non-adaptive, the protocol is non-adaptive as well.

Proof. Apply Theorem 8.3 with $\lambda(n) = n$, $p(\lambda) = \lambda$, and empty advice. □

The one-bit arguments from witness encryption and iO in Sections 7.1 and 7.1 have $c_0 = 1/2$ and preserve the number of reduction queries. Thus the corollary applies to their uniform multi-query reductions under the corresponding target-hardness hypothesis, with T bounding the composed soundness reduction.

9 From MIPs to Witness and Public-Key Encryption

We construct witness encryption from MIPs with short prover replies and then obtain public-key encryption. Encryption security reduces to the PRG; NIWI witness indistinguishability is used only for correctness.

9.1 Witness Encryption

Following [LMP25], we first construct a weak witness encryption scheme (Theorem 9.1), and then amplify it to full witness encryption using Theorem 4.10.

Ingredients.

- A perfectly complete, statistically sound NIWI given by algorithms $(\mathcal{G}_{\mathsf{NIWI}}, \mathcal{P}_{\mathsf{NIWI}}, \mathcal{V}_{\mathsf{NIWI}})$ (Definition 4.23), with $\mathsf{poly}(T, 2^\ell, \lambda)$-secure witness indistinguishability. We write ν_{NIWI} for its statistical soundness error.

- For the NP relation $\mathcal{R}'$ under consideration, nondecreasing bounds $T(N) \geq N$ and $\ell(N) \geq 1$, and an MIP $(\mathcal{P}_1, \ldots, \mathcal{P}_{k'}, \mathcal{V})$. On statements of length N, there are $k' \leq \mathsf{poly}(T(N))$ provers, the verifier runs in time $\mathsf{poly}(T(N))$, and the honest provers run in time $\mathsf{poly}(N)$. The verifier contacts at most $k(N) \leq T(N)$ provers, each at most once. Each prover receives only the common input, its identity, and its question, and sends at most $\ell(N) \leq \mathsf{poly}(N)$ bits. Besides the statement, the verifier receives an advice string z of length at most $S(N) \leq T(N)$, and there is a fixed advice string z_N for each length N. The honest provers ignore the advice. The protocol has strong local correctness (Definition 8.1) with error $\epsilon_c(N)$ for *every* advice string z, and statistical soundness error $\epsilon_s(N)$ when $z = z_N$, where

$$1 - \epsilon_c(N) - \epsilon_s(N) \geq \frac{1}{\mathsf{poly}(N)}.$$

  For uniform reductions, z_N is the empty string.

Without loss of generality, the verifier contacts exactly k distinct provers.

The relation $\mathcal{R}_{\mathsf{NIWI}}$. Let $\mathcal{R}$ be an NP relation with language L and witness bound $m(n)$, and fix a PRG $\mathsf{PRG} : \{0,1\}^\kappa \to \{0,1\}^{2\kappa}$. For $x \in \{0,1\}^n$ and $y \in \{0,1\}^{2\kappa}$, set $m' = 1 + \max\{m(|x|), \kappa\}$ and define $\mathcal{R}_{\vee}$ using the following equal-length witnesses for (x, y):

$$\begin{aligned} &w' = 0 \parallel w \parallel 0^{m'-1-|w|} && \text{for } (x,w) \in \mathcal{R}, \\ &w' = 1 \parallel s \parallel 0^{m'-1-\kappa} && \text{for } s \in \{0,1\}^\kappa \text{ with } y = \mathsf{PRG}(s). \end{aligned} \tag{14}$$

For $X = (\mathsf{crs}, x, y)$, define the relation $\mathcal{R}_{\mathsf{NIWI}}$ by

$$\big((\mathsf{crs}, x, y), \pi\big) \in \mathcal{R}_{\mathsf{NIWI}} \quad \Longleftrightarrow \quad \mathcal{V}_{\mathsf{NIWI}}(\mathsf{crs}, (x, y), \pi) = 1.$$

Take an MIP as above for $\mathcal{R}_{\mathsf{NIWI}}$ and pad the CRS so that $N = |X|$ depends only on (n, λ). Write $z_{n,\lambda} = z_N$ for its advice.

Parameters. For encryption security parameter λ, choose a $\kappa = \kappa(n, \lambda) \leq \mathsf{poly}(n, \lambda)$ with $n \leq \mathsf{poly}(\kappa)$ and $\kappa \geq \log \lambda$. For a fixed constant $d \geq 1$, the compiled statements and NIWI proofs have length at most κ^d. Write $T = T(\kappa^d)$, $\ell = \ell(\kappa^d)$, and $\Sigma = \{0,1\}^{\leq \ell}$.

Let α, β be functions of (n, λ) with $\beta \leq \alpha \leq \mathsf{poly}(T, n, \lambda)$, and set

$$\eta = \Theta\left(\frac{1}{|\Sigma|^2 \alpha(n,\lambda)}\right), \qquad M = \Theta\left(\frac{\lambda + n + \log|\Sigma| + \log \alpha}{\eta^2}\right),$$

with sufficiently small and large constants, respectively.

The scheme.

Encryption $\mathsf{Enc}'(1^\lambda, x)$, for $|x| = n$, with advice $z = z_{n,\lambda}$:

1. Sample $s \leftarrow \{0,1\}^\kappa$ and $\mathsf{crs} \leftarrow \mathcal{G}_{\mathsf{NIWI}}(1^\kappa, 1^{|x|+2\kappa}, 1^{m'})$. Set $y = \mathsf{PRG}(s)$, $X = (\mathsf{crs}, x, y)$, and $s' = 1 \parallel s \parallel 0^{m'-1-\kappa}$.

2. Run $\mathcal{V}(X, z)$. When it queries prover i on u, sample a fresh $\pi \leftarrow \mathcal{P}_{\mathsf{NIWI}}(\mathsf{crs}, (x, y), s')$ and reply with $\mathcal{P}_i(X, u, \pi)$. Sample $j \leftarrow [k]$, and let (i, u) be the j-th prover contacted and its question.

3. Independently sample $\pi_1, \ldots, \pi_M \leftarrow \mathcal{P}_{\mathsf{NIWI}}(\mathsf{crs}, (x, y), s')$.

   Let $\hat{h}$ be the histogram of $\mathcal{P}_i(X, u, \pi_t)$ for $t \in [M]$, using fresh prover coins each time.

4. Sample $v \leftarrow \{0, \ldots, M-1\}$. Output $\mathsf{ct} = (x, \mathsf{crs}, y, i, u, v)$ and $\mu = \mathsf{CorSamp}(\hat{h}, v)$.

Here x is the statement, crs the NIWI reference string, y the PRG output, (i, u) the selected prover and its question, and v the common random index used by $\mathsf{CorSamp}$.

Decryption $\mathsf{Dec}'(\mathsf{ct}, w)$, for $\mathsf{ct} = (x, \mathsf{crs}, y, i, u, v)$:

1. Set $X = (\mathsf{crs}, x, y)$ and $w' = 0 \parallel w \parallel 0^{m'-1-|w|}$.
2. Independently sample $\pi_1, \ldots, \pi_M \leftarrow \mathcal{P}_{\mathsf{NIWI}}(\mathsf{crs}, (x, y), w')$.
   Let $\hat{h}'$ be the histogram of $\mathcal{P}_i(X, u, \pi_t)$ for $t \in [M]$, using fresh prover coins each time.
3. Output $\mathsf{CorSamp}(\hat{h}', v)$.

We assume $\mathsf{poly}(T, 2^\ell, \lambda)$-security of the PRG and NIWI.
The choices of η and M give

$$e^{-2M\eta^2} \le \frac{e^{-(\lambda+n)}}{|\Sigma|^2 \alpha}, \qquad M \le \mathsf{poly}(T, 2^\ell, \lambda).$$

For the analysis, write $\mathcal{D}_s(X, i, u)$ and $\mathcal{D}_w(X, i, u)$ for the single-reply distributions sampled in Step 3 and decryption, respectively. Both algorithms run in time $\mathsf{poly}(T, M, \kappa) = \mathsf{poly}(T, 2^\ell, \lambda)$, while the ciphertext and sampled string have length $\mathsf{poly}(T, n, \log \lambda)$. Since the verifier may choose each question based on earlier replies, Step 2 runs the full execution. The advice only affects which provers are contacted and which questions they receive. Since the honest provers ignore it, $\mathcal{D}_s(X, i, u)$ and $\mathcal{D}_w(X, i, u)$ do not depend on it, and decryption does not use it.

Theorem 9.1. *Suppose that $\mathcal{R}_{\mathsf{NIWI}}$ admits the MIPs specified under Ingredients. Let*

$$\beta(n, \lambda) = \Theta\left(\max_N \frac{k(N)}{1 - \epsilon_c(N) - \epsilon_s(N)}\right),$$

where the maximum is over the statement lengths N at security parameter $\kappa(n, \lambda)$ and the constant is sufficiently large. Assume $\mathsf{poly}(T, 2^\ell, \lambda)$-security of the PRG and NIWI, and choose α as in Parameters. Then $(\mathsf{Enc}', \mathsf{Dec}')$ runs in time $\mathsf{poly}(T, 2^\ell, \lambda)$, and has correctness $1 - 1/\alpha(n, \lambda)$ and hardness $1/\beta(n, \lambda)$ on statements of length n, where $\beta(n, \lambda) \le \mathsf{poly}(T, \kappa)$. Correctness holds for every advice string of length at most $S(N)$ given to Enc', and hardness holds for the advice $z_{n,\lambda}$. Hardness holds against adversaries of size $\mathsf{poly}(T, 2^\ell, \lambda)$.

Moreover, from any adversary $\mathcal{A}$ violating hardness on a statement $x \notin L$, we obtain, in a black-box way, a distinguisher for PRG at seed length κ with advantage at least $(1 - \epsilon_c - \epsilon_s)/2$, making at most k calls to $\mathcal{A}$ and of size $|\mathcal{A}| \cdot k + \mathsf{poly}(T, n, \lambda)$, with x and $z_{n,\lambda}$ hardwired.

Proof. We prove correctness and hardness separately.

Claim 9.2. $(\mathsf{Enc}', \mathsf{Dec}')$ *has correctness $1 - 1/\alpha(n, \lambda)$.*

Proof. Fix a sufficiently large λ and $(x, w) \in \mathcal{R}$ with $|x| = n$. Fix also s and the advice string z given to Enc'; we prove the bound for every such choice. The argument uses no property of the MIP. One encryption and decryption use $k + 2M$ NIWI proofs for the statement (x, y) under the same CRS: k in Step 2, M in Step 3, and M in Dec'. The first $k + M$ use the encoded seed s', and the last M use the encoded witness w'. Let Fail be the event that the decrypted string differs from μ.

Switching the decryptor's witness. For $t \in \{0, \dots, M\}$, let Exp_t be the honest execution, except that the first t proofs in Dec' use s' instead of w'. Thus, Exp_0 is the honest execution. Given crs and a single proof π, the following circuit $\mathcal{B}_t$ computes Fail in Exp_t or Exp_{t+1}, depending on whether π uses w' or s'. It generates all other proofs itself, which is possible since the CRS is public and s', w', z are hardwired, uses π as the $(t+1)$-st proof in Dec', and runs the rest of the encryption and decryption. Its size is $\mathsf{poly}(T, M, \kappa) = \mathsf{poly}(T, 2^\ell, \lambda)$. Since $\alpha M \le \mathsf{poly}(T, 2^\ell, \lambda)$, the assumed NIWI security gives

$$\big|\Pr[\mathsf{Fail} \text{ in } \mathsf{Exp}_t] - \Pr[\mathsf{Fail} \text{ in } \mathsf{Exp}_{t+1}]\big| \le \nu_{\mathsf{WI}} = o(1/(\alpha M)).$$

Summing over t, the failure probabilities in Exp_0 and Exp_M differ by at most $M\nu_{\mathsf{WI}}$.

Both histograms estimate the same distribution in Exp_M. Condition on crs and (i, u). In Exp_M, both $\hat{h}$ and $\hat{h}'$ are histograms of M independent samples from $\mathcal{D}_s(X, i, u)$, since each sample uses a fresh proof and fresh coins of $\mathcal{P}_i$. By Lemmas 4.39 and 4.38, and since the norm is at most 2,

$$\Pr[\mathsf{Fail} \text{ in } \mathsf{Exp}_M] \le |\Sigma| \cdot \mathbb{E}\left[\left\|\frac{\hat{h}}{M} - \frac{\hat{h}'}{M}\right\|_1\right] \le 2|\Sigma|^2\eta + 8e^{-(\lambda+n)}/\alpha.$$

Combining the two bounds,

$$\Pr[\mathsf{Fail} \text{ in } \mathsf{Exp}_0] \le O\big(|\Sigma|^2\eta + e^{-(\lambda+n)}/\alpha\big) + M\nu_{\mathsf{WI}}.$$

Here $M\nu_{\mathsf{WI}} = o(1/\alpha)$, and the exponential term is also $o(1/\alpha)$. By the choice of η with a sufficiently small constant, the first term is at most $1/(2\alpha(n, \lambda))$. Thus, the error is at most $1/\alpha(n, \lambda)$ for sufficiently large λ. □

Remark 9.3. *The proof of Claim 9.2 uses witness indistinguishability only against the circuits $\mathcal{B}_t$. These run the honest encryption and decryption algorithms $(\mathsf{Enc}', \mathsf{Dec}')$ constructed above, with x, w, s, and z hardwired; in particular, they run the MIP verifier with honest provers, which, in our applications, simulates the underlying soundness reduction with the honest prover. Thus, as for security against $\mathcal{R} \circ \mathcal{P}$, it suffices that every circuit in this class has advantage at most ν_{WI} in the witness-indistinguishability game of the NIWI, with $M\nu_{\mathsf{WI}} = o(1/\alpha)$.*

Claim 9.4. $(\mathsf{Enc}', \mathsf{Dec}')$ *has hardness $1/\beta(n, \lambda)$, where $\beta(n, \lambda) \le \mathsf{poly}(T, \kappa)$, with the adversary size and black-box reduction stated in the theorem.*

Proof. Suppose that, for an adversary $\mathcal{A}$ of size $\mathsf{poly}(T, 2^\ell, \lambda)$, there are infinitely many λ with a statement $x \notin L$ of length n such that

$$\Pr\left[\ \mathcal{A}(\mathsf{ct}) \neq \mu \ : \ (\mathsf{ct}, \mu) \leftarrow \mathsf{Enc}'(1^\lambda, x)\ \right] < \frac{1}{\beta}.$$

Fix such a λ and x, and drop the subscript of $\mathcal{A}$.

Cheating provers. For a statement $X = (\mathsf{crs}, x, y)$, define prover $\mathcal{P}_i^*$ as follows: on question u, sample $v \leftarrow \{0, \dots, M-1\}$ and return $\mathcal{A}(x, \mathsf{crs}, y, i, u, v)$. Each $\mathcal{P}_i^*$ uses only X, its identity, its question, and independent private coins, so $(\mathcal{P}_1^*, \dots, \mathcal{P}_{k'}^*)$ is a valid tuple of noncommunicating provers.

For $b \in \{0, 1\}$, let E_b be the experiment that samples $\mathsf{crs} \leftarrow \mathcal{G}_{\mathsf{NIWI}}(1^\kappa, 1^{|x|+2\kappa}, 1^{m'})$ and $y \in \{0,1\}^{2\kappa}$, where $y = \mathsf{PRG}(s)$ for $s \leftarrow \{0,1\}^\kappa$ if $b = 0$, and y is uniform if $b = 1$. It then runs $\mathcal{V}(X, z)$ with the cheating provers, for $X = (\mathsf{crs}, x, y)$ and $z = z_{n,\lambda}$. Let p_b be its acceptance probability.

The cheating provers convince the verifier in E_0. We show that

$$p_0 \geq 1 - \epsilon_c(N) - \frac{k}{\beta}. \tag{15}$$

In E_0, let H be the execution of $\mathcal{V}(X, z)$ in Step 2 of encryption, and let H^* be its execution with the cheating provers. By perfect completeness of the NIWI, every proof sampled in H is a valid witness for X under $\mathcal{R}_{\mathsf{NIWI}}$. Strong local correctness of the MIP, which holds for every advice string, therefore gives

$$\Pr[H \text{ accepts}] \geq 1 - \epsilon_c(N). \tag{16}$$

It suffices to bound the statistical distance between the two executions.

Fix (s, crs). For $j \in [k]$, let Z_j consist of the verifier's tape ρ, the first $j-1$ contacted provers with their questions and replies, and the j-th contacted prover i_j with its question u_j. Let d_j be the statistical distance between the distributions of (Z_j, σ_j) in H and H^*, where σ_j is the j-th reply, and set $d_0 = 0$. The verifier's output is a function of (Z_k, σ_k), so it suffices to bound d_k.

Each prover is contacted at most once. Thus, conditional on (i_j, u_j), the reply σ_j in H has distribution $\mathcal{D}_s(X, i_j, u_j)$ and is independent of the rest of Z_j. In H^*, it is distributed according to the distribution $\mathcal{A}_{i_j,u_j}$ of $\mathcal{A}(x, \mathsf{crs}, y, i_j, u_j, v)$ for uniform v, again independently of the rest of Z_j. Moreover, Z_{j+1} is obtained from (Z_j, σ_j) by the same verifier computation in both executions. Therefore,

$$d_j \leq d_{j-1} + \mathop{\mathbb{E}}_{Z_j \sim H}\Big[\mathrm{SD}\big(\mathcal{D}_s(X, i_j, u_j), \mathcal{A}_{i_j,u_j}\big)\Big].$$

Here, d_{j-1} accounts for the difference in the distribution of (i_j, u_j) between the two executions, which arises when the verifier is adaptive.

To bound the summand, fix (i, u), and let $\hat{h}$ be the histogram from Step 3 of Enc'. Since $\hat{h}$ is the histogram of M independent samples from $\mathcal{D}_s(X, i, u)$, and $\mathsf{CorSamp}(\hat{h}, v)$ is distributed according to $\hat{h}/M$ for uniform v, the string $\mathsf{CorSamp}(\hat{h}, v)$ is distributed exactly according to $\mathcal{D}_s(X, i, u)$. Let

$$f(i, u) = \Pr\big[\mathcal{A}(x, \mathsf{crs}, y, i, u, v) \neq \mathsf{CorSamp}(\hat{h}, v)\big],$$

over $\hat{h}$, v, and the coins of $\mathcal{A}$. Coupling the two outputs through the same v,

$$\mathrm{SD}\big(\mathcal{D}_s(X, i, u), \mathcal{A}_{i,u}\big) \leq f(i, u).$$

Summing over j, and using that Enc' chooses $j \leftarrow [k]$ uniformly, with $(i, u) = (i_j, u_j)$ distributed as in H,

$$d_k \leq \sum_{j=1}^{k} \mathop{\mathbb{E}}_{Z_j \sim H}\big[f(i_j, u_j)\big] = k \cdot \Pr\big[\mathcal{A}(\mathsf{ct}) \neq \mu \mid s, \mathsf{crs}\big].$$

Averaging over (s, crs) and using the assumption on $\mathcal{A}$ gives (15).

The verifier rejects in E_1. In E_1, the string y lies outside the image of PRG, except with probability $2^{-\kappa}$. Since $x \notin L$, in this case $(x, y) \notin L_{\mathcal{R}_\vee}$, so statistical soundness of the NIWI gives $X \in L_{\mathcal{R}_{\mathsf{NIWI}}}$ with probability at most $\nu_{\mathsf{NIWI}}(\kappa)$ over the CRS. On the remaining CRSs, the verifier runs on (X, z_N) with $X \notin L_{\mathcal{R}_{\mathsf{NIWI}}}$, so soundness of the MIP for the advice z_N bounds the acceptance probability of the cheating provers by $\epsilon_s(N)$. This is the only step that uses the correctness of the advice. Thus,

$$p_1 \leq \epsilon_s(N) + 2^{-\kappa} + \nu_{\mathsf{NIWI}}(\kappa).$$

A PRG distinguisher. The distinguisher $\mathcal{B}$, with x and $z = z_{n,\lambda}$ hardwired, takes y as input, samples crs, runs $\mathcal{V}(X, z)$ with the cheating provers, and outputs the verifier's decision. Its advantage in distinguishing $\mathsf{PRG}(s)$ from a uniform string is $p_0 - p_1$.

Conclusion. Combining the bounds,

$$p_0 - p_1 \geq (1 - \epsilon_c - \epsilon_s) - \frac{k}{\beta} - 2^{-\kappa} - \nu_{\mathsf{NIWI}}(\kappa).$$

By the choice of β with a sufficiently large constant, $k/\beta \leq (1 - \epsilon_c - \epsilon_s)/4$. By the parameter choice above, $(1 - \epsilon_c - \epsilon_s) \geq 1/\operatorname{poly}(\kappa)$, while $2^{-\kappa}$ and $\nu_{\mathsf{NIWI}}(\kappa)$ are negligible in κ. Thus, the advantage of $\mathcal{B}$ is at least $(1 - \epsilon_c - \epsilon_s)/2 \geq 1/\operatorname{poly}(\kappa)$ for sufficiently large λ.

On the other hand, $\mathcal{B}$ has size $|\mathcal{A}| \cdot k + \operatorname{poly}(T, n, \lambda) = \operatorname{poly}(T, 2^\ell, \lambda)$. Since $T \geq \kappa$, the assumed $\operatorname{poly}(T, 2^\ell, \lambda)$-security of PRG makes its advantage smaller than $(1 - \epsilon_c - \epsilon_s)/2$, a contradiction. Finally, $\beta(n, \lambda) \leq \operatorname{poly}(T, \kappa)$, since $k \leq T$ and $1/(1 - \epsilon_c - \epsilon_s) \leq \operatorname{poly}(\kappa)$. □

Claims 9.2 and 9.4 together prove the theorem. □

Amplification.

Theorem 9.5. *Let $\mathcal{R}$ be an* NP *relation. Suppose that $\mathcal{R}_{\mathsf{NIWI}}$ admits an MIP with reply bound ℓ and time bound T, as specified under Ingredients. Assume $\operatorname{poly}(T, 2^\ell, \lambda)$-security of the PRG and NIWI. Then there is a witness encryption scheme for $\mathcal{R}$ satisfying the correctness and security conditions of Definition 4.8, with encryption and decryption time*

$$\operatorname{poly}(T, 2^\ell, \lambda)$$

on statements of length n, with $T = T(\kappa^d)$ and $\ell = \ell(\kappa^d)$ as above. Its security has a black-box reduction to the PRG with overhead $\operatorname{poly}(T, 2^\ell, \lambda)$. In particular, when $T, \lambda \leq 2^{O(\ell)}$, both the running time and the required security are $\operatorname{poly}(2^\ell)$. If the MIP uses advice, encryption takes the advice $z_{n,\lambda}$, of length at most T, as a non-uniform input; correctness holds for every advice string, and security holds for $z_{n,\lambda}$.

Proof. Choose $\beta(n, \lambda)$ as in Theorem 9.1, the amplification parameter $B = n + \lambda + \lceil T \rceil$, and $\alpha(n, \lambda) = 4B\beta(n, \lambda)$, so that $\alpha, \beta \leq \operatorname{poly}(T, n, \lambda)$. Then $\alpha, \beta \leq \operatorname{poly}(B)$, as required even when T is superpolynomial. Apply the amplification of Theorem 4.10 with amplification parameter B, passing the same advice $z_{n,\lambda}$ to every call of Enc'. The amplification uses correctness and hardness of the weak scheme only for the fixed statement x of length n, so it applies unchanged. It makes $\operatorname{poly}(B)$ calls to the weak scheme, so the running time remains $\operatorname{poly}(T, 2^\ell, \lambda)$ and correctness error is negligible in $n + \lambda$. A polynomial-size distinguisher yields a weak-scheme adversary of size $\operatorname{poly}(T, 2^\ell, \lambda)$, which is covered by Theorem 9.1. Composing the two reductions gives the claimed black-box reduction. □

For the subexponential instantiation, suppose that the PRG is ϵ-subexponentially secure and the NIWI has ϵ-subexponential witness indistinguishability, for some $0 < \epsilon \leq 1$. Suppose also that $\mathcal{R}_{\mathsf{NIWI}}$ admits the MIPs above with reply bound $N^{\epsilon'}$ for every constant $\epsilon' > 0$. Fix a constant $\delta > 0$, and set

$$\kappa = \left\lceil (\log \lambda + n^\delta)^{2/\epsilon} \right\rceil, \qquad \epsilon' = \frac{\epsilon}{2d}.$$

Then $n \leq \kappa^{\epsilon/(2\delta)}$ and each reply has at most $\ell = \lceil \kappa^{\epsilon/2} \rceil = O(\log \lambda + n^\delta)$ bits. If

$$\log T = O(\kappa^{\epsilon/2}), \tag{17}$$

then $\operatorname{poly}(T, 2^\ell, \lambda) = 2^{O(\kappa^{\epsilon/2})}$, so subexponential security suffices. Theorem 9.5 therefore gives encryption and decryption time $\operatorname{poly}(T, \lambda) 2^{O(n^\delta)}$. Condition (17) holds, for instance, when T is polynomial.

9.2 Public-Key Encryption

We obtain public-key encryption using the construction of [GGSW13]. Let $\mathcal{R}_{\mathsf{PRG}} = \{(y, s) : y = \mathsf{PRG}(s)\}$, so that a statement $y \in \{0,1\}^{2\kappa}$ has witnesses in $\{0,1\}^{\kappa}$.

Corollary 9.6. *Suppose that the PRG and NIWI satisfy the subexponential-security hypotheses following Theorem 9.5. Consider the relation with statements (crs, y_1, y_2) and witnesses given by accepting NIWI proofs of "y_1 or y_2 is a PRG image". Assume that this relation admits the MIPs specified under Ingredients, with reply bound $N^{\epsilon'}$ for every constant $\epsilon' > 0$ and time bound T satisfying (17). Then there is a public-key encryption scheme, with encryption and decryption time $\mathsf{poly}(T, \lambda)$, whose security has a black-box reduction to the subexponential security of the PRG. If the MIPs use advice, its encryption algorithm is non-uniform, with one advice string of length at most T per security parameter; key generation and decryption are uniform.*

Proof. Apply the subexponential parameter choice following Theorem 9.5 to $\mathcal{R}_{\mathsf{PRG}}$ with $\delta = \epsilon/3$, and let $(\mathsf{Enc}, \mathsf{Dec})$ be the resulting witness encryption scheme. Set $\kappa_0(\lambda) = \lceil (\log \lambda)^{2/\epsilon} \rceil$. Key generation samples $s \leftarrow \{0,1\}^{\kappa_0(\lambda)}$ and outputs $\mathsf{pk} = \mathsf{PRG}(s)$ and $\mathsf{sk} = s$. Encryption of b outputs $\mathsf{Enc}(1^\lambda, \mathsf{pk}, b)$, and decryption outputs $\mathsf{Dec}(\mathsf{ct}, \mathsf{sk})$.

The public key has length $n = 2\kappa_0(\lambda)$, so $n^\delta = O((\log \lambda)^{2/3})$ and encryption and decryption run in time $\mathsf{poly}(T, \lambda)$. Since n is determined by λ, the advice $z_{n,\lambda}$ of the witness encryption is a single string per security parameter, which we hardwire into the encryption algorithm; key generation and decryption do not use it. Correctness follows from correctness of the witness encryption.

For security, every polynomial-size distinguisher has advantage at most $\lambda^{-\log \lambda}$ against PRG at seed length $\kappa_0(\lambda)$. Thus, replacing $\mathsf{pk} = \mathsf{PRG}(s)$ by a uniformly random $y \in \{0,1\}^{2\kappa_0(\lambda)}$ changes the adversary's advantage by a negligible amount. This y is not a PRG image except with negligible probability $2^{-\kappa_0(\lambda)}$, and security of witness encryption on these false statements hides the encrypted bit. Composing the reductions gives the claim. □

9.3 Implications

We now combine Theorem 9.5 with the MIPs of Section 8. For witness encryption and iO, the one-bit laconic arguments of Section 7.1 give MIPs in which each contacted prover sends a single bit. The construction then runs in time $\mathsf{poly}(T)$ and needs only $\mathsf{poly}(T)$-security of the PRG and the NIWI, so a polynomial-time multi-query reduction yields polynomial-time witness encryption and public-key encryption whose security reduces to a one-way function. For SNARGs, the replies have length $N^{\epsilon'}$, and we use the subexponential instantiation following Theorem 9.5.

Witness encryption.

Corollary 9.7. *Suppose WE for NP admits a multi-query black-box security reduction to a game-based assumption $(\mathcal{C}, c)$. Suppose the reduction uses only length-dependent advice. Let $(\mathcal{R}, F)$ be the soundness reduction for the one-bit argument in Section 7.1, with running time bounded by $T_{\mathcal{R}}(\lambda, N)$, and set $T(N) = T_{\mathcal{R}}(N, N)$. Take T nondecreasing with $T(N) \geq N$.*

Fix a length-doubling PRG and a NIWI. For each NP relation, there is a constant $d \geq 1$ such that the following holds. Additionally, we assume that:

- *The NIWI is perfectly complete and statistically sound. We require only class-dependent $\mathsf{poly}(T(\kappa^d))$-secure witness indistinguishability against the circuits $\mathcal{B}_t$ that run honest encryption and decryption (Remark 9.3).*
- *The PRG has $\mathsf{poly}(T(\kappa^d))$-security at primitive security parameter κ.*

- *$(\mathcal{C}, c)$ is secure against the honest-prover reduction $\mathcal{R} \circ \mathcal{P}$ at the reduction's target-parameter map, for every polynomial bound on the source index length. We require only this class-dependent security (Definition 6.1).*

Then there is a non-uniform witness encryption scheme for the relation with encryption and decryption time

$$\mathsf{poly}\big(T((n+\lambda)^d), n, \lambda\big).$$

Its security reduces to the PRG via a reduction that treats the encryption adversary as a black box, with overhead polynomial in the same quantities.

In particular, when T is polynomial, polynomial witness indistinguishability against the circuits $\mathcal{B}_t$ and polynomially secure one-way functions suffice. The scheme then runs in polynomial time, and its security reduces to the one-way function.

Proof. The construction in Section 7.1 gives a one-bit laconic argument with soundness baseline $1/2$, preserving the reduction and its advice. Apply Theorem 8.3 with $\lambda(N) = N$ to obtain MIPs with one-bit replies and verifier time $\mathsf{poly}(T(N))$; the uniform case is Corollary 8.5. Set $\kappa = n + \lambda$, and choose a constant d bounding the compiled statement and NIWI proof lengths by κ^d. Since $T(\kappa^d) \geq \kappa$ and $\ell = 1$, the assumed security suffices in Theorem 9.5, with the NIWI restriction of Remark 9.3. This gives the claimed WE time and reduction overhead.

When T is polynomial, the required security is polynomial. A polynomially secure one-way function gives such a PRG [HILL99]; composing the reductions proves the final claim. □

Corollary 9.8. *Under the hypotheses of Corollary 9.7 for the relation $\mathcal{R}_{\mathsf{PRG}}$ and a sufficiently large constant d, there is a public-key encryption scheme with encryption and decryption time*

$$\mathsf{poly}\big(T(\lambda^d), \lambda\big).$$

Its security has a black-box reduction to the PRG with overhead polynomial in the same quantities. When T is polynomial, polynomial witness indistinguishability against the class in Remark 9.3 and polynomially secure one-way functions suffice; the scheme runs in polynomial time, and its security reduces to the one-way function.

Proof. Apply Corollary 9.7 to $\mathcal{R}_{\mathsf{PRG}}$. Generate $\mathsf{pk} = (1^\lambda, \mathsf{PRG}(s))$ and $\mathsf{sk} = s$ for $s \leftarrow \{0,1\}^\lambda$, and encrypt using the resulting WE. Here $n = 2\lambda$. For a sufficiently large d, encryption, decryption, and the PRG distinguisher run in time $\mathsf{poly}(T(\lambda^d), \lambda)$. The assumed PRG security at seed length λ therefore permits replacing the PRG output in the public key by a uniform string, which is outside the PRG image except with probability $2^{-\lambda}$. WE security then hides the encrypted bit. The polynomial-time specialization follows as in Corollary 9.7. □

Obfuscation. Corollaries 9.7 and 9.8 also apply starting from iO for circuits. The standard construction of WE from iO [GGH+13] has a one-query black-box security reduction (Section 7.1). Composing it with the assumed iO reduction preserves the number of oracle queries and length-dependent advice. The time bound $T_{\mathcal{R}}$ and the class-dependent security hypotheses are those of the composed reduction.

SNARGs. Combining Corollary 9.6 with Corollary 8.4 and [HILL99], using the NIWI security restriction of Remark 9.3, we obtain the following.

Corollary 9.9. *Suppose a non-adaptively sound SNARG for NP admits a polynomial-time, multi-query black-box soundness reduction $(\mathcal{R}, F)$ to a game-based assumption $(\mathcal{C}, c)$. Suppose the reduction uses only length-dependent advice. Additionally, we assume that:*

- *Perfectly complete, statistically sound NIWIs exist. We require only class-dependent, subexponential witness indistinguishability against the circuits $\mathcal{B}_t$ that run honest encryption and decryption* (*Remark 9.3*).
- *Subexponentially secure one-way functions exist.*
- $(\mathcal{C}, c)$ *is secure against the honest-prover reduction* $\mathcal{R} \circ \mathcal{P}$ *at the reduction's target-parameter map, for every polynomial bound on the statement length. We require only this class-dependent security* (*Definition 6.1*).

Then there is a public-key encryption scheme whose security has a black-box reduction to the subexponential security of the one-way function. With length-dependent advice, its encryption algorithm is non-uniform.

For length-dependent advice, Theorem 8.3 supplies the required MIPs with $z_N = F(1^{\lambda(N)}, 1^N)$: strong local correctness holds for every advice string, and soundness holds for z_N. Encryption uses one advice string per (n, λ) for WE and per λ for PKE; decryption and key generation are uniform. The uniform case is the special case of empty advice. Statement-dependent advice is not covered, since encryption would need to compute the correct advice for each freshly sampled $X = (\mathsf{crs}, x, y)$.

NIWI witness indistinguishability is used only for correctness. Encryption security reduces to the PRG. When the PRG is constructed from a one-way function, composing the reductions gives security from that function. The construction uses the PRG description in the OR relation, so the reduction is black-box in the encryption adversary, while the construction is not fully black-box in the one-way function.

10 AI Methodology

Technical content. All core mathematical ideas in this paper were developed by the authors, with the following exceptions.

- For Theorem 6.6, the authors initially attempted an approach based on Chernoff bounds; Claude Opus 5.5 suggested the more elegant argument via Carathéodory's theorem.
- The one-query lower bound for adaptive soundness follows from our laconic interactive proof construction via Corollary 6.9. The extension to multi-query reductions in Appendix A was developed by Codex GPT 6 Astra, with interactive guidance from the authors.
- During a local "STOC review" of the paper by GPT 6 Astra, it pointed out that our general theorems yield stronger conclusions for iO and WE (relative to SNARGs) since they yield laconic arguments with one-bit prover messages.

Writing. The paper was written with assistance from Codex GPT 6 Astra, Claude Opus 5.5 and Fable 5.1. We also used them to search the literature.

The authors have checked all AI-assisted content, including proofs and citations, and take full responsibility for the paper.

11 References

[AGGM06] Adi Akavia, Oded Goldreich, Shafi Goldwasser, and Dana Moshkovitz. On basing one-way functions on NP-hardness. In Jon M. Kleinberg, editor, *38th ACM STOC*, pages 701–710, Seattle, WA, USA, May 21–23, 2006. ACM Press. 14

[AJ15] Prabhanjan Ananth and Abhishek Jain. Indistinguishability obfuscation from compact functional encryption. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 308–326, Santa Barbara, CA, USA, August 16–20, 2015. Springer Berlin Heidelberg, Germany. 1, 3, 4

[Bab85] László Babai. Trading group theory for randomness. In *17th ACM STOC*, pages 421–429, Providence, RI, USA, May 6–8, 1985. ACM Press. 11

[BB15] Andrej Bogdanov and Christina Brzuska. On basing size-verifiable one-way functions on NP-hardness. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part I*, volume 9014 of *LNCS*, pages 1–6, Warsaw, Poland, March 23–25, 2015. Springer Berlin Heidelberg, Germany. 14

[BBK+23] Zvika Brakerski, Maya Farber Brodsky, Yael Tauman Kalai, Alex Lombardi, and Omer Paneth. SNARGs for monotone policy batch NP. In Helena Handschuh and Anna Lysyanskaya, editors, *CRYPTO 2023, Part II*, volume 14082 of *LNCS*, pages 252–283, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland. 3, 76

[BDRV18] Itay Berman, Akshay Degwekar, Ron D. Rothblum, and Prashant Nalini Vasudevan. From laconic zero-knowledge to public-key cryptography - extended abstract. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 674–697, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Cham, Switzerland. 6, 11

[Bel02] Mihir Bellare. A note on negligible functions. *Journal of Cryptology*, 15(4):271–284, September 2002. 9, 38

[BGI+01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In Joe Kilian, editor, *CRYPTO 2001*, volume 2139 of *LNCS*, pages 1–18, Santa Barbara, CA, USA, August 19–23, 2001. Springer Berlin Heidelberg, Germany. 1, 20

[BHZ87] Ravi B. Boppana, Johan Håstad, and Stathis Zachos. Does co-NP have short interactive proofs? *Information Processing Letters*, 25(2):127–132, 1987. 3, 13

[BIOW20] Ohad Barta, Yuval Ishai, Rafail Ostrovsky, and David J. Wu. On succinct arguments and witness encryption from groups. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 776–806, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Cham, Switzerland. 2, 16

[BJK+18] Zvika Brakerski, Aayush Jain, Ilan Komargodski, Alain Passelègue, and Daniel Wichs. Non-trivial witness encryption and null-iO from standard assumptions. In Dario Catalano and Roberto De Prisco, editors, *SCN 18*, volume 11035 of *LNCS*, pages 425–441, Amalfi, Italy, September 5–7, 2018. Springer, Cham, Switzerland. 5

[BKK+18] Saikrishna Badrinarayanan, Yael Tauman Kalai, Dakshita Khurana, Amit Sahai, and Daniel Wichs. Succinct delegation for low-space non-deterministic computation. In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *50th ACM STOC*, pages 709–721, Los Angeles, CA, USA, June 25–29, 2018. ACM Press. 76, 77

[BM82] Manuel Blum and Silvio Micali. How to generate cryptographically strong sequences of pseudo random bits. In *23rd FOCS*, pages 112–117, Chicago, Illinois, November 3–5, 1982. IEEE Computer Society Press. 22

[BM88] László Babai and Shlomo Moran. Arthur–Merlin games: A randomized proof system, and a hierarchy of complexity classes. *Journal of Computer and System Sciences*, 36(2):254–276, 1988. 26

[BR94] Mihir Bellare and John Rompel. Randomness-efficient oblivious sampling. In *35th FOCS*, pages 276–287, Santa Fe, NM, USA, November 20–22, 1994. IEEE Computer Society Press. 28

[BT03] Andrej Bogdanov and Luca Trevisan. On worst-case to average-case reductions for NP problems. In *44th FOCS*, pages 308–317, Cambridge, MA, USA, October 11–14, 2003. IEEE Computer Society Press. 14

[BV15] Nir Bitansky and Vinod Vaikuntanathan. Indistinguishability obfuscation from functional encryption. In Venkatesan Guruswami, editor, *56th FOCS*, pages 171–190, Berkeley, CA, USA, October 17–20, 2015. IEEE Computer Society Press. 1, 3, 4

[CCL18] Yi-Hsiu Chen, Kai-Min Chung, and Jyun-Jie Liao. On the complexity of simulating auxiliary input. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 371–390, Tel Aviv, Israel, April 29 – May 3, 2018. Springer, Cham, Switzerland. 2

[CGKS23] Matteo Campanelli, Chaya Ganesh, Hamidreza Khoshakhlagh, and Janno Siim. Impossibilities in succinct arguments: Black-box extraction and more. In Nadia El Mrabet, Luca De Feo, and Sylvain Duquesne, editors, *AFRICACRYPT 23*, volume 14064 of *LNCS*, pages 465–489, Sousse, Tunisia, July 19–21, 2023. Springer, Cham, Switzerland. 2

[CHS05] Ran Canetti, Shai Halevi, and Michael Steiner. Hardness amplification of weakly verifiable puzzles. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 17–33, Cambridge, MA, USA, February 10–12, 2005. Springer Berlin Heidelberg, Germany. 19

[CJ25] Liyan Chen and Zhengzhong Jin. On the impossibility of SNARGs with short CRS : (or: Revisiting gentry-wichs barrier in the non-adaptive setting). In *66th FOCS*, pages 1741–1760, Sydney, Australia, December 14–17, 2025. IEEE Computer Society Press. 1, 2, 3, 6, 11, 14, 17, 20, 48, 49, 77

[CJJ21] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. Non-interactive batch arguments for NP from standard assumptions. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 394–423, Virtual Event, August 16–20, 2021. Springer, Cham, Switzerland. 3

[CJJ22] Arka Rai Choudhuri, Abhishek Jain, and Zhengzhong Jin. SNARGs for $\mathcal{P}$ from LWE. In *62nd FOCS*, pages 68–79, Denver, CO, USA, February 7–10, 2022. IEEE Computer Society Press. 3, 76

[CLP13] Ran Canetti, Huijia Lin, and Rafael Pass. From unprovability to environmentally friendly protocols. In *54th FOCS*, pages 70–79, Berkeley, CA, USA, October 26–29, 2013. IEEE Computer Society Press. 14, 15

[CLP15] Kai-Min Chung, Edward Lui, and Rafael Pass. From weak to strong zero-knowledge and applications. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part I*,

volume 9014 of *LNCS*, pages 66–92, Warsaw, Poland, March 23–25, 2015. Springer Berlin Heidelberg, Germany. 2

[COA20] Francesco Cosentino, Harald Oberhauser, and Alessandro Abate. A randomized algorithm to reduce the support of discrete measures. In *Advances in Neural Information Processing Systems*, 2020. 11, 42

[CPW20] Suvradip Chakraborty, Manoj Prabhakaran, and Daniel Wichs. Witness maps and applications. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part I*, volume 12110 of *LNCS*, pages 220–246, Edinburgh, UK, May 4–7, 2020. Springer, Cham, Switzerland. 4

[CPW23] Suvradip Chakraborty, Manoj Prabhakaran, and Daniel Wichs. A map of witness maps: New definitions and connections. In Alexandra Boldyreva and Vladimir Kolesnikov, editors, *PKC 2023, Part II*, volume 13941 of *LNCS*, pages 635–662, Atlanta, GA, USA, May 7–10, 2023. Springer, Cham, Switzerland. 4

[DHK+26] Lalita Devadas, Samuel B. Hopkins, Yael Tauman Kalai, Pravesh K. Kothari, Alex Lombardi, and Surya Mathialagan. SNARGs for NP and non-signaling PCPs, revisited. In *58th ACM STOC*. ACM Press, 2026. Full version: Cryptology ePrint Archive, Paper 2026/006. 1, 76

[DMP26] Nico Döttling, Giulio Malavolta, and Omer Paneth. Batch arguments with optimal communication. Cryptology ePrint Archive, Paper 2026/003, 2026. 77

[FF93] Joan Feigenbaum and Lance Fortnow. Random-self-reducibility of complete sets. *SIAM J. Comput.*, 22(5):994–1005, 1993. 14

[FLS90] Uriel Feige, Dror Lapidot, and Adi Shamir. Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In *31st FOCS*, pages 308–317, St. Louis, MO, USA, October 22–24, 1990. IEEE Computer Society Press. 11

[GGH+13] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49, Berkeley, CA, USA, October 26–29, 2013. IEEE Computer Society Press. 1, 2, 7, 20, 46, 62

[GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 467–476, Palo Alto, CA, USA, June 1–4, 2013. ACM Press. 1, 4, 13, 18, 61

[GL89] Oded Goldreich and Leonid A. Levin. A hard-core predicate for all one-way functions. In *21st ACM STOC*, pages 25–32, Seattle, WA, USA, May 15–17, 1989. ACM Press. 19

[GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, 1984. 24

[GMM17a] Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammed. Lower bounds on obfuscation from all-or-nothing encryption primitives. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 661–695, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Cham, Switzerland. 14

[GMM17b] Sanjam Garg, Mohammad Mahmoody, and Ameer Mohammed. When does functional encryption imply obfuscation? In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 82–115, Baltimore, MD, USA, November 12–15, 2017. Springer, Cham, Switzerland. 14

[GOS06] Jens Groth, Rafail Ostrovsky, and Amit Sahai. Perfect non-interactive zero knowledge for NP. In Serge Vaudenay, editor, *EUROCRYPT 2006*, volume 4004 of *LNCS*, pages 339–358, St. Petersburg, Russia, May 28 – June 1, 2006. Springer Berlin Heidelberg, Germany. 23

[GOS12] Jens Groth, Rafail Ostrovsky, and Amit Sahai. New techniques for noninteractive zero-knowledge. *J. ACM*, 59(3), June 2012. 23

[GPSZ17] Sanjam Garg, Omkant Pandey, Akshayaram Srinivasan, and Mark Zhandry. Breaking the sub-exponential barrier in obfustopia. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part III*, volume 10212 of *LNCS*, pages 156–181, Paris, France, April 30 – May 4, 2017. Springer, Cham, Switzerland. 1, 3

[GRS23] Oded Goldreich, Guy N. Rothblum, and Tal Skverer. On interactive proofs of proximity with proof-oblivious queries. In Yael Tauman Kalai, editor, *ITCS 2023*, volume 251, pages 59:1–59:16, Cambridge, MA, USA, January 10–13, 2023. LIPIcs. 26

[GS86] Shafi Goldwasser and Michael Sipser. Private coins versus public coins in interactive proof systems. In *18th ACM STOC*, pages 59–68, Berkeley, CA, USA, May 28–30, 1986. ACM Press. 11, 26

[GSWW22] Rachit Garg, Kristin Sheridan, Brent Waters, and David J. Wu. Fully succinct batch arguments for np from indistinguishability obfuscation. In Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022, Part I*, volume 13747 of *LNCS*, pages 526–555, Chicago, IL, USA, November 7–10, 2022. Springer, Cham, Switzerland. 77

[GVW02] Oded Goldreich, Salil Vadhan, and Avi Wigderson. On interactive proofs with a laconic prover. *Computational Complexity*, 11(1–2):1–53, 2002. 6, 10, 14, 25, 26, 27, 41

[GW11] Craig Gentry and Daniel Wichs. Separating succinct non-interactive arguments from all falsifiable assumptions. In Lance Fortnow and Salil P. Vadhan, editors, *43rd ACM STOC*, pages 99–108, San Jose, CA, USA, June 6–8, 2011. ACM Press. 1, 2, 5, 11, 13, 15, 17, 45, 71, 72, 75

[HILL99] Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999. 23, 62

[HJLM26] Yao-Ching Hsieh, Abhishek Jain, Jiatu Li, and Surya Mathialagan. SNARGs for NP from unprovability of mathematical theorems. In *58th ACM STOC*. ACM Press, 2026. Full version: Cryptology ePrint Archive, Paper 2026/1180. 1, 76

[HK14] Thomas Holenstein and Robin Künzler. A protocol for generating random elements with their probabilities. In *Computing and Combinatorics – COCOON 2014*, volume 8591 of *LNCS*, pages 191–202. Springer, 2014. Full version: https://arxiv.org/abs/1312.2483. 25, 26

[HMX10] Iftach Haitner, Mohammad Mahmoody, and David Xiao. A new sampling protocol and applications to basing cryptographic primitives on the hardness of NP. In *25th IEEE CCC*, pages 76–87, 2010. 14

[Hoe63] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963. 28

[IR89] Russell Impagliazzo and Steven Rudich. Limits on the provable consequences of one-way permutations. In *21st ACM STOC*, pages 44–61, Seattle, WA, USA, May 15–17, 1989. ACM Press. 5, 13

[Jin26] Zhengzhong Jin. Low-norm Nullstellensatz hypothesis for the AND code is false. Cryptology ePrint Archive, Paper 2026/1087, 2026. 76

[JJ22] Abhishek Jain and Zhengzhong Jin. Indistinguishability obfuscation via mathematical proofs of equivalence. In *63rd FOCS*, pages 1023–1034, Denver, CO, USA, October 31 – November 3, 2022. IEEE Computer Society Press. 1, 3, 77

[JJMP25] Abhishek Jain, Zhengzhong Jin, Surya Mathialagan, and Omer Paneth. On succinct obfuscation via propositional proofs. In *66th FOCS*, pages 1703–1740, Sydney, Australia, December 14–17, 2025. IEEE Computer Society Press. 77

[JKLM25] Zhengzhong Jin, Yael Tauman Kalai, Alex Lombardi, and Surya Mathialagan. Universal SNARGs for NP from proofs of correctness. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 933–943, Prague, Czechia, June 23–27, 2025. ACM Press. 1, 3, 9, 16, 76, 77

[JKLV24] Zhengzhong Jin, Yael Kalai, Alex Lombardi, and Vinod Vaikuntanathan. SNARGs under LWE via propositional proofs. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 1750–1757, Vancouver, BC, Canada, June 24–28, 2024. ACM Press. 3, 9, 76, 77

[JLP26] Zhengzhong Jin, Mingqi Lu, and Bo Peng. SNARKs from LWE via non-black-box reductions. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing (STOC 2026)*, pages 1812–1823, 2026. 76

[JLS21] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from well-founded assumptions. In Samir Khuller and Virginia Vassilevska Williams, editors, *53rd ACM STOC*, pages 60–73, Virtual Event, Italy, June 21–25, 2021. ACM Press. 1, 77

[JLS22] Aayush Jain, Huijia Lin, and Amit Sahai. Indistinguishability obfuscation from LPN over $\mathbb{F}_p$, DLIN, and PRGs in NC^0. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 670–699, Trondheim, Norway, May 30 – June 3, 2022. Springer, Cham, Switzerland. 1

[JP14] Dimitar Jetchev and Krzysztof Pietrzak. How to fake auxiliary input. In Yehuda Lindell, editor, *TCC 2014*, volume 8349 of *LNCS*, pages 566–590, San Diego, CA, USA, February 24–26, 2014. Springer Berlin Heidelberg, Germany. 2

[Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *24th ACM STOC*, pages 723–732, Victoria, BC, Canada, May 4–6, 1992. ACM Press. 1

[KLMR18] Yuan Kang, Chengyu Lin, Tal Malkin, and Mariana Raykova. Obfuscation from polynomial hardness: Beyond decomposable obfuscation. In Dario Catalano and Roberto De Prisco, editors, *SCN 18*, volume 11035 of *LNCS*, pages 407–424, Amalfi, Italy, September 5–7, 2018. Springer, Cham, Switzerland. 1, 77

[KVZ21] Yael Tauman Kalai, Vinod Vaikuntanathan, and Rachel Yun Zhang. Somewhere statistical soundness, post-quantum security, and SNARGs. In Kobbi Nissim and Brent Waters, editors, *TCC 2021, Part I*, volume 13042 of *LNCS*, pages 330–368, Raleigh, NC, USA, November 8–11, 2021. Springer, Cham, Switzerland. 76

[LMP25] Yanyi Liu, Noam Mazor, and Rafael Pass. On witness encryption and laconic zero-knowledge arguments. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 429–461, Santa Barbara, CA, USA, August 17–21, 2025. Springer, Cham, Switzerland. 6, 11, 13, 19, 29, 55

[LPWW20] Benoît Libert, Alain Passelègue, Hoeteck Wee, and David J. Wu. New constructions of statistical NIZKs: Dual-mode DV-NIZKs and more. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 410–441, Zagreb, Croatia, May 10–14, 2020. Springer, Cham, Switzerland. 23

[LY94] Richard J. Lipton and Neal E. Young. Simple strategies for large zero-sum games with applications to complexity theory. In *26th ACM STOC*, pages 734–740, Montréal, Québec, Canada, May 23–25, 1994. ACM Press. 13, 72

[LZ17] Qipeng Liu and Mark Zhandry. Decomposable obfuscation: A framework for building applications of obfuscation from polynomial hardness. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017, Part I*, volume 10677 of *LNCS*, pages 138–169, Baltimore, MD, USA, November 12–15, 2017. Springer, Cham, Switzerland. 1, 3, 77

[MDS25] Yaohua Ma, Chenxin Dai, and Elaine Shi. Quasi-linear indistinguishability obfuscation via mathematical proofs of equivalence and applications. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part III*, volume 15603 of *LNCS*, pages 157–186, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland. 1, 77

[Mic94] Silvio Micali. CS proofs (extended abstracts). In *35th FOCS*, pages 436–453, Santa Fe, NM, USA, November 20–22, 1994. IEEE Computer Society Press. 1

[MMN+16] Mohammad Mahmoody, Ameer Mohammed, Soheil Nematihaji, Rafael Pass, and abhi shelat. Lower bounds on assumptions behind indistinguishability obfuscation. In Eyal Kushilevitz and Tal Malkin, editors, *TCC 2016-A, Part I*, volume 9562 of *LNCS*, pages 49–66, Tel Aviv, Israel, January 10–13, 2016. Springer Berlin Heidelberg, Germany. 14

[MPV24] Surya Mathialagan, Spencer Peters, and Vinod Vaikuntanathan. Adaptively sound zero-knowledge SNARKs for UP. In Leonid Reyzin and Douglas Stebila, editors, *CRYPTO 2024, Part X*, volume 14929 of *LNCS*, pages 38–71, Santa Barbara, CA, USA, August 18–22, 2024. Springer, Cham, Switzerland. 2

[Nao03] Moni Naor. On cryptographic assumptions and challenges (invited talk). In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 96–109, Santa Barbara, CA, USA, August 17–21, 2003. Springer Berlin Heidelberg, Germany. 1

[PS19] Chris Peikert and Sina Shiehian. Noninteractive zero knowledge for NP from (plain) learning with errors. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part I*, volume 11692 of *LNCS*, pages 89–114, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland. 23

[RTV04] Omer Reingold, Luca Trevisan, and Salil P. Vadhan. Notions of reducibility between cryptographic primitives. In Moni Naor, editor, *TCC 2004*, volume 2951 of *LNCS*, pages 1–20, Cambridge, MA, USA, February 19–21, 2004. Springer Berlin Heidelberg, Germany. 5

[Rud97] Steven Rudich. Super-bits, demi-bits, and NP/qpoly-natural proofs. In *Randomization and Approximation Techniques in Computer Science (RANDOM 1997)*, volume 1269 of *Lecture Notes in Computer Science*, pages 85–93. Springer, 1997. 1

[RV10] Guy N. Rothblum and Salil P. Vadhan. Are PCPs inherent in efficient arguments? *Computational Complexity*, 19(2):265–304, 2010. Preliminary version in CCC 2009. 3, 14

[RVV24] Seyoon Ragavan, Neekon Vafa, and Vinod Vaikuntanathan. Indistinguishability obfuscation from bilinear maps and LPN variants. In Elette Boyle and Mohammad Mahmoody, editors, *TCC 2024, Part IV*, volume 15367 of *LNCS*, pages 3–36, Milan, Italy, December 2–6, 2024. Springer, Cham, Switzerland. 1

[Sio58] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958. 11, 50

[SW14] Amit Sahai and Brent Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In David B. Shmoys, editor, *46th ACM STOC*, pages 475–484, New York, NY, USA, May 31 – June 3, 2014. ACM Press. 1, 3, 76, 77

[Tha23] Justin Thaler. Proofs, arguments, and zero-knowledge, 2023. Online manuscript, July 18, 2023. 27

[Vad12] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1–3):1–336, 2012. 29

[vN28] John von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928. 24

[Wil16] R. Ryan Williams. Strong ETH breaks with Merlin and Arthur: Short non-interactive proofs of batch evaluation. In *31st Conference on Computational Complexity (CCC 2016)*, volume 50 of *LIPIcs*, pages 2:1–2:17. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2016. 4

[Wu18] David J. Wu. *Lattice-Based Non-Interactive Argument Systems*. PhD thesis, Stanford University, 2018. 45

[WW24] Brent Waters and David J. Wu. Adaptively-sound succinct arguments for NP from indistinguishability obfuscation. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *56th ACM STOC*, pages 387–398, Vancouver, BC, Canada, June 24–28, 2024. ACM Press. 76

[WW25] Brent Waters and David J. Wu. A pure indistinguishability obfuscation approach to adaptively-sound SNARGs for NP. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part VII*, volume 16006 of *LNCS*, pages 292–326, Santa Barbara, CA, USA, August 17–21, 2025. Springer, Cham, Switzerland. 76

[WWW25] Brent Waters, Hoeteck Wee, and David J. Wu. New techniques for preimage sampling: Improved NIZKs and more from LWE. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part IV*, volume 15604 of *LNCS*, pages 3–33, Madrid, Spain, May 4–8, 2025. Springer, Cham, Switzerland. 23

[Yao82] Andrew Chi-Chih Yao. Theory and applications of trapdoor functions (extended abstract). In *23rd FOCS*, pages 80–91, Chicago, Illinois, November 3–5, 1982. IEEE Computer Society Press. 22

A Multi-Query Lower Bounds for Adaptive SNARGs

In this section, we prove Theorem 6.10. We first begin by giving an overview of the main idea behind the result in Section A.1. Then, we recap relevant tools from Gentry-Wichs [GW11] in Section A.2. We then give the proof in Section A.3, and conclude with the implication for adaptive SNARGs in Section A.4.

A.1 Technical Overview

We show that the argument of Gentry and Wichs [GW11] only needs a worst-case hypothesis: if an adaptively sound SNARG for L has a k-query black-box reduction to a secure falsifiable assumption, then L has circuits of size $2^{O(\ell_{\mathcal{P}})}\,\mathrm{poly}(T)$ (Theorem 6.10). We follow the same two steps as our one-query results.

1. We use the reduction to construct a statistically sound *oracle test* (Theorem A.4). Instead of receiving a statement x, the test interacts with an oracle that answers each query with a statement of its choice together with a proof. It accepts oracles that give true statements with honest proofs, and rejects oracles that only give false statements. The assumption is used only for completeness, and the reduction only for soundness.

2. We show that any such test gives small circuits for L (Theorem A.5). This step uses no cryptography. If L were hard, the lemma of Gentry and Wichs would give an oracle that only outputs false statements, yet is indistinguishable from an honest oracle. The test accepts the honest oracle and rejects this one, so it would distinguish the two.

The oracle chooses a fresh statement on every query. If every query used the same statement x, the replies would be correlated through x, and the Gentry–Wichs lemma, which handles $\ell_{\mathcal{P}}$ bits at a time, would no longer apply. Adaptive soundness gives exactly this independence. We replace the hard subset-membership problem of [GW11] by hard distributions obtained from worst-case hardness via a standard minimax argument, and since these distributions need not be efficiently samplable, completeness uses a fixed list of instance–witness pairs.

A.2 Tools and Definitions from Prior Work

Two distributions D_Y, D_N are (s,ε)-*indistinguishable* if every Boolean circuit C with at most s gates satisfies

$$\left| \Pr_{x \leftarrow D_Y}[C(x)=1] - \Pr_{x \leftarrow D_N}[C(x)=1] \right| \leq \varepsilon.$$

An *augmentation* of D_Y is a joint distribution of pairs (x,π) whose x-marginal is D_Y. We use the following lemma of Gentry and Wichs.

Lemma A.1 ([GW11, Lemma 3.1])**.** *There is a fixed polynomial Q_{GW} with the following property. Let $\ell \geq 0$ be an integer, $s_0 \geq 1$, and $0 < \varepsilon < 1/2$, and let D_Y, D_N be (s,ε)-indistinguishable, where*

$$s \geq (s_0+1)\, Q_{\mathrm{GW}}(2^\ell, 1/\varepsilon).$$

Then for every augmentation (x,π) of D_Y with $\pi \in \{0,1\}^\ell$, there is an augmentation $(x,\widetilde{\pi})$ of D_N with $\widetilde{\pi} \in \{0,1\}^\ell$ such that the two joint distributions are $(s_0, 2\varepsilon)$-indistinguishable.

Note that none of these distributions need to be efficiently samplable, and π may be arbitrarily correlated with x, for example through a witness [GW11, Remarks on Lemma 3.1].

Let $\mathsf{size}_L(n)$ denote the minimum number of gates of a circuit deciding L on all n-bit strings. Gentry and Wichs obtain D_Y, D_N from a hard subset-membership problem. We instead obtain them from worst-case hardness, using a standard minimax argument (see, e.g., the proof of [LY94, Theorem 6]).

Lemma A.2. *Let $s \geq 1$ and $0 < \varepsilon < 1/2$. If*

$$\mathsf{size}_L(n) > \Theta\Big(\frac{n}{\varepsilon^2}\Big) \cdot (s+1) \tag{18}$$

for a sufficiently large hidden constant, then there are (s,ε)-indistinguishable distributions D_Y on $L \cap \{0,1\}^n$ and D_N on $\{0,1\}^n \setminus L$.

Proof. Suppose the conclusion fails, and let $\mathcal{S}$ be the set of circuits on n inputs with at most $s+1$ gates. Since $\mathcal{S}$ contains the negation of every circuit with at most s gates, for every distribution D_Y on $L \cap \{0,1\}^n$ and D_N on $\{0,1\}^n \setminus L$, there is a circuit $C \in \mathcal{S}$ with

$$\Pr_{y \leftarrow D_Y}[C(y)=1] - \Pr_{x \leftarrow D_N}[C(x)=1] > \varepsilon.$$

Consider the finite zero-sum game in which one player chooses a pair $(y,x) \in (L \cap \{0,1\}^n) \times (\{0,1\}^n \setminus L)$, the other chooses $C \in \mathcal{S}$, and the second player receives payoff $C(y) - C(x)$. The expected payoff of a mixed strategy on pairs depends only on its two marginals, so every mixed strategy of the first player has a response with expected payoff greater than ε. By the minimax theorem (Theorem 4.27), there is a distribution μ on $\mathcal{S}$ such that

$$\mathbb{E}_{C \leftarrow \mu}[C(y)] - \mathbb{E}_{C \leftarrow \mu}[C(x)] \geq \varepsilon \qquad \text{for every } y \in L \cap \{0,1\}^n \text{ and } x \in \{0,1\}^n \setminus L.$$

In particular, there is a threshold θ such that $\mathbb{E}_\mu[C(y)] \geq \theta + \varepsilon/2$ for every $y \in L \cap \{0,1\}^n$, and $\mathbb{E}_\mu[C(x)] \leq \theta - \varepsilon/2$ for every $x \in \{0,1\}^n \setminus L$.

Now sample $C_1, \ldots, C_t \leftarrow \mu$ independently, for $t = \Theta(n/\varepsilon^2)$ with a sufficiently large constant. For each fixed input x, by Hoeffding's inequality (Lemma 4.33), the average $\frac{1}{t}\sum_i C_i(x)$ lies on the wrong side of θ with probability at most $e^{-t\varepsilon^2/2} < 2^{-n}$. By a union bound over all $x \in \{0,1\}^n$, there is a fixed choice of $C_1, \ldots, C_t$ for which the comparison of this average with θ decides L on every n-bit input. Comparing a sum of t bits with a threshold takes $O(t)$ gates, so L is decided on $\{0,1\}^n$ by a circuit of size $t(s+1) + O(t)$, which contradicts (18). □

A.3 From Multi-Query Reductions to Small Circuits

Fix λ, n and a proof system $\Pi = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ for an NP relation R_L. Proofs together with a failure symbol are encoded in $\ell_{\mathcal{P}} + 1$ bits.

Definition A.3. *A* stateless oracle $\mathcal{O}$ *answers each query* crs *with a pair* $(x, \pi) \leftarrow \mathcal{O}(\mathsf{crs})$*, using fresh coins on each call. We say* $\mathcal{O}$ *is a* false-statement oracle *if every pair it outputs satisfies* $|x| = n$ *and* $x \notin L$*. For a distribution* D_Y *on* $L \cap \{0,1\}^n$ *and witnesses* $w(x)$ *for each* x *in its support, the* honest oracle $\mathcal{H}_{D_Y, w}$ *samples* $x \leftarrow D_Y$ *and returns* $(x, \mathcal{P}(\mathsf{crs}, x, w(x)))$*.*

Neither D_Y nor w needs to be efficiently computable. An *oracle test* for Π at length n is an oracle algorithm $\mathcal{V}'$ with completeness error ϵ_c and soundness error ϵ_s if

$$\Pr\left[\mathcal{V}'^{\mathcal{H}_{D_Y, w}} \text{ accepts}\right] \geq 1 - \epsilon_c \qquad \text{and} \qquad \Pr\left[\mathcal{V}'^{\mathcal{O}} \text{ accepts}\right] \leq \epsilon_s$$

for every D_Y and w as above, and every possibly unbounded false-statement oracle $\mathcal{O}$.

Without loss of generality, the time bound $T(\lambda, n)$ of the reduction also bounds all lengths and running times below.

Theorem A.4. *Fix a nondecreasing polynomial* p *and a security-parameter function* $\lambda = \lambda(n)$ *such that* $n \leq p(\lambda(n))$ *for every* n*. Suppose* L *has an adaptively sound SNARG* $\Pi = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ *with proof length* $\ell_{\mathcal{P}}(\lambda, n)$*, whose soundness admits a non-uniform,* $k(\lambda, n)$*-query,* $T(\lambda, n)$*-time black-box reduction* $(\mathcal{R}, F)$ *with adaptive queries to a game-based assumption* $(\mathcal{C}, c)$*.*

Then, assuming that $(\mathcal{C}, c)$ *is secure against* $\mathcal{R} \circ \mathcal{P}$ *at the reduction's target-parameter map* f *(Definition 6.1), there exist completeness and soundness error bounds* ϵ_c, ϵ_s *satisfying*

$$1 - \epsilon_c(n) - \epsilon_s(n) \geq \frac{1}{\mathsf{poly}(\lambda(n), n)}$$

such that, for all sufficiently large n*, there is an oracle test* $\mathcal{V}'$ *for* Π *at length* n *and security parameter* $\lambda(n)$ *with errors* $\epsilon_c(n), \epsilon_s(n)$*. Given its random coins, which include advice* $z \leftarrow F(1^{\lambda(n)}, 1^n)$*, the oracle test runs in time* $O(T(\lambda(n), n))$ *and makes at most* $k(\lambda(n), n)$ *queries.*

Proof. Apply Lemma 6.2, with $\Gamma(\lambda, n) = \lambda$, to the adaptive soundness games of Definition 4.7, whose public index is 1^n. Thus there is an inverse-polynomial bound $\delta(\lambda, n) > 0$ such that, for all sufficiently large λ and every $n \leq p(\lambda)$, every possibly unbounded stateless oracle $\mathcal{O}$ with adaptive soundness advantage at least $1/2$ satisfies

$$\Pr\left[\mathcal{R}^{\mathcal{O}}(1^\lambda, 1^n, F(1^\lambda, 1^n)) \text{ wins } \mathcal{C}(1^{\lambda'})\right] \geq c + \delta(\lambda, n), \tag{19}$$

where $\lambda' = f(1^\lambda, n)$. The probability includes the independently sampled advice $F(1^\lambda, 1^n)$. We may take δ to be at most $(1 - c)/2$.

Choose a constant $\beta \in (2c/(1 + 2c), 1)$, and put $\tau = 1 - c(1 - \beta)$. We use the protocol from Theorem 6.8, with the following changes: there is no common input, and the oracle chooses the statement.

On input $(1^\lambda, 1^n)$, the verifier $\mathcal{V}'$ has oracle access to $\mathcal{O}$.

1. With probability β, the verifier samples $(\mathsf{crs}, \mathsf{vk}) \leftarrow \mathcal{G}(1^\lambda, 1^n)$ and queries $(x, \pi) \leftarrow \mathcal{O}(\mathsf{crs})$. It accepts exactly when $|x| = n$ and $\mathcal{V}(\mathsf{vk}, x, \pi) = 1$.
2. Otherwise, it samples $z \leftarrow F(1^\lambda, 1^n)$ and simulates $\mathcal{R}(1^\lambda, 1^n, z)$ interacting with $\mathcal{C}(1^{\lambda'})$

using independent coins, forwarding each oracle query crs of $\mathcal{R}$ to $\mathcal{O}$ and returning its reply. It accepts exactly when $\mathcal{C}$ does not output win.

Completeness. For every k-tuple $\mathbf{t} = ((x_j, w_j))_{j=1}^{k}$ of statement–witness pairs with $|x_j| = n$, and z with $|z| \le T(\lambda, n)$, the assumption adversary $\mathcal{B}_{\mathbf{t},z}$ with $\mathbf{t}, z$ hardwired can run $\mathcal{R}(1^\lambda, 1^n, z)$ and answer its jth query crs_j with $(x_j, \mathcal{P}(\mathsf{crs}_j, x_j, w_j))$. This adversary is $(\mathcal{R} \circ \mathcal{P})(1^n, z, \mathbf{t})$ (Definition 6.1), so security against $\mathcal{R} \circ \mathcal{P}$ gives a common negligible bound $\nu(\lambda)$ on its advantage at λ', which we increase to cover the original completeness error. Now fix D_Y and w. Since $\mathcal{H}_{D_Y,w}$ samples its statement independently of the query, $\mathcal{R}^{\mathcal{H}_{D_Y,w}}(1^\lambda, 1^n, F(1^\lambda, 1^n))$ is distributed as $\mathcal{B}_{\mathbf{t},z}$ for independent $x_1, \dots, x_k \leftarrow D_Y$, $w_j = w(x_j)$, and $z \leftarrow F(1^\lambda, 1^n)$. Averaging over $\mathbf{t}$ and z, the mixed verifier therefore accepts $\mathcal{H}_{D_Y,w}$ with probability at least $\tau - \nu(\lambda)$ for every D_Y and w.

Statistical soundness. Fix n and a possibly unbounded false-statement oracle $\mathcal{O}$. Let α_{cheat} be its adaptive soundness advantage. Since every statement output by $\mathcal{O}$ is false, this is exactly the acceptance probability in the first branch. The mixed acceptance probability is

$$\beta\alpha_{\mathrm{cheat}} + (1-\beta)\left(1 - \Pr\left[\mathcal{R}^{\mathcal{O}}(1^\lambda, 1^n, F(1^\lambda, 1^n)) \text{ wins } \mathcal{C}(1^{\lambda'})\right]\right).$$

If $\alpha_{\mathrm{cheat}} < 1/2$, this is at most $\beta/2 + 1 - \beta$. Otherwise, apply (19). The same calculation as in Theorem 6.3 bounds acceptance by $\tau - \xi(\lambda, n)$, where

$$\xi(\lambda, n) = \min\left\{\frac{\beta}{2} - (1-\beta)c,\ (1-\beta)\delta(\lambda, n)\right\}.$$

Choose a cutoff so that $\nu(\lambda) \le \xi(\lambda, n)/2$ for all larger λ and all $n \le p(\lambda)$. Set

$$\epsilon_c(n) = 1 - \tau + \frac{\xi(\lambda(n), n)}{2}, \qquad \epsilon_s(n) = \tau - \xi(\lambda(n), n).$$

These bounds have gap $\xi/2$. Since $n \le p(\lambda(n))$, the security parameter $\lambda(n)$ exceeds the cutoff for all sufficiently large n. Given its coins, including z, the mixed verifier runs in time $O(T(\lambda, n))$ and makes at most $k(\lambda, n)$ queries. □

We now show that an oracle test with short proofs implies small circuits for L. This step does not use any cryptographic assumption.

Theorem A.5. *Let $\mathcal{V}'$ be an oracle test for Π at length n with errors ϵ_c, ϵ_s satisfying $\Delta = 1 - \epsilon_c - \epsilon_s > 0$. Suppose that, given its random coins, $\mathcal{V}'$ runs in time T and makes at most $k \le T$ queries, and that T also bounds the running time of $\mathcal{P}$ and the witness lengths. Then*

$$\mathsf{size}_L(n) \le 2^{O(\ell_{\mathcal{P}})}\,\mathsf{poly}(T, 1/\Delta).$$

Proof. Choose $\varepsilon = \Theta(\Delta/k)$ with a sufficiently small constant so that $2k\varepsilon < \Delta$, let S_2 be a sufficiently large fixed polynomial, and set

$$s = (S_2(T) + 1)\, Q_{\mathrm{GW}}(2^{\ell_{\mathcal{P}}+1}, 1/\varepsilon).$$

Suppose for contradiction that $\mathsf{size}_L(n)$ satisfies (18). By Lemma A.2, there are (s, ε)-indistinguishable distributions D_Y on $L \cap \{0,1\}^n$ and D_N on $\{0,1\}^n \setminus L$. Fix witnesses $w(x)$ for each x in the support of D_Y, and write $\mathcal{H} = \mathcal{H}_{D_Y,w}$.

A cheating oracle. For each crs, the reply distribution of $\mathcal{H}(\mathsf{crs})$ is an augmentation of D_Y. Following [GW11, proof of Lemma 4.1], we apply Lemma A.1 with $\ell = \ell_{\mathcal{P}} + 1$ and $s_0 = S_2(T)$. This gives an augmentation of D_N that is $(S_2(T), 2\varepsilon)$-indistinguishable from $\mathcal{H}(\mathsf{crs})$. Let $\mathcal{A}(\mathsf{crs})$ sample this augmentation with fresh coins on each call. In particular, every statement output by $\mathcal{A}$ is false, so $\mathcal{A}$ is a false-statement oracle. By completeness and soundness of $\mathcal{V}'$, we therefore have that

$$\Pr[\mathcal{V}'^{\mathcal{H}} \text{ accepts}] - \Pr[\mathcal{V}'^{\mathcal{A}} \text{ accepts}] \geq 1 - \epsilon_c - \epsilon_s = \Delta > 2k\varepsilon. \tag{20}$$

Replacing one reply at a time. For $0 \leq i \leq k$, consider the execution of $\mathcal{V}'$ in which the first i queries are answered by $\mathcal{A}$ and the rest by $\mathcal{H}$. By (20), the acceptance probabilities of some two consecutive executions differ by more than 2ε. These executions differ only in the reply to one query. By averaging, we can fix the coins of $\mathcal{V}'$, the replies to the earlier queries, and the pairs $(x, w(x))$ and prover coins used for the later replies, while retaining this gap. This fixes the query crs at the changed position.

We can now define a distinguisher for $\mathcal{A}(\mathsf{crs})$ and $\mathcal{H}(\mathsf{crs})$. Given a reply to crs, it runs the rest of $\mathcal{V}'$, computing each later reply from its fixed pair. This is possible even though the later queries depend on the given reply, since the fixed pairs do not. The distinguisher is a circuit of size at most $S_2(T)$ with advantage greater than 2ε, which contradicts the choice of $\mathcal{A}$. Hence

$$\mathsf{size}_L(n) \leq \Theta\Big(\frac{n}{\varepsilon^2}\Big) \cdot (s+1) = 2^{O(\ell_{\mathcal{P}})}\, \mathsf{poly}(T, 1/\Delta). \qquad \square$$

A.4 Implication for Adaptive SNARGs

We now prove Theorem 6.10, which we restate here.

Theorem 6.10 (Restated)**.** *Fix a nondecreasing polynomial p and a security-parameter function $\lambda = \lambda(n)$ such that $n \leq p(\lambda(n))$ for every n. Suppose L has an adaptively sound SNARG $\Pi = (\mathcal{G}, \mathcal{P}, \mathcal{V})$ with proof length $\ell_{\mathcal{P}}(\lambda, n)$, whose soundness admits a non-uniform, $k(\lambda, n)$-query, $T(\lambda, n)$-time black-box reduction $(\mathcal{R}, F)$ with adaptive queries to a game-based assumption $(\mathcal{C}, c)$. Assume that $(\mathcal{C}, c)$ is secure against $\mathcal{R} \circ \mathcal{P}$ at the reduction's target-parameter map f (Definition 6.1). Then, writing $\lambda = \lambda(n)$,*

$$L \in \mathsf{SIZE}\Big[2^{O(\ell_{\mathcal{P}}(\lambda,n))} \cdot \mathsf{poly}\big(T(\lambda, n)\big)\Big].$$

Proof. The hypotheses are exactly those of Theorem A.4. For all sufficiently large n, Theorem A.4 gives an oracle test $\mathcal{V}'$ for Π at length n that, given its random coins, runs in time $O(T(\lambda, n))$, makes at most $k(\lambda, n) \leq T(\lambda, n)$ queries, and has gap $\Delta = 1 - \epsilon_c - \epsilon_s$ with $1/\Delta \leq \mathsf{poly}(\lambda, n) \leq \mathsf{poly}(T(\lambda, n))$. Applying Theorem A.5 to $\mathcal{V}'$ gives the claimed bound. $\square$

The following form of Corollary 7.8 makes the dependence on the security-parameter exponent explicit. The advice $F(1^\lambda, 1^n)$ may have any description complexity and is sampled independently of the cheating prover.

Corollary A.6. *Let $\alpha, d > 0$ be constants, and choose $\lambda(n) = \Theta(n^d)$. Let p be a polynomial with $n \leq p(\lambda(n))$ for every n. Suppose $L \in \mathsf{NP}$ has an adaptively sound SNARG with proof length $\ell_{\mathcal{P}}(\lambda, n) = O(\lambda^\alpha)(n + q(n))^{o(1)}$, whose adaptive soundness admits a non-uniform, $k(\lambda, n)$-query, $T(\lambda, n)$-time black-box reduction with adaptive queries to a game-based assumption $(\mathcal{C}, c)$. Assume that $(\mathcal{C}, c)$ is T-hard at the reduction's target-parameter map f, for statement lengths $n \leq p(\lambda)$. Then*

$$L \in \mathsf{SIZE}\Big[2^{n^{\alpha d + o(1)}} \cdot \mathsf{poly}\big(T(\lambda(n), n)\big)\Big].$$

In particular, for every constant $\epsilon > 0$, choosing $d < \epsilon/\alpha$ gives circuits of size $2^{n^\epsilon} \cdot \mathsf{poly}(T(\lambda(n), n))$.

Proof. Apply Theorem 6.10. The assumed hardness is exactly the hypothesis of Theorem A.4. Since q is polynomial, the proof-length bound becomes $\ell_{\mathcal{P}}(\lambda(n), n) = O(n^{\alpha d}) \cdot n^{o(1)} = n^{\alpha d + o(1)}$, and the claimed bound follows. If $\alpha d < \epsilon$, then $n^{\alpha d + o(1)} \leq n^{\epsilon}$ for all sufficiently large n. □

For example, if $T(\lambda, n) = 2^{O(n^{\gamma})}\,\mathsf{poly}(\lambda)$ for a constant $\gamma > 0$, then Corollary A.6 gives circuits of size $2^{n^{\epsilon}} \cdot 2^{O(n^{\gamma})}$, assuming security against adversaries of size $2^{O(n^{\gamma})} = 2^{O(\lambda^{\gamma/d})}$ at the target parameter $\lambda' = f(1^{\lambda}, n)$, where $\lambda = \Theta(n^d)$.

When $k = 1$, Theorem A.5 covers the setting of Theorem 5.2. There, the circuit is obtained constructively by enumerating the $2^{\ell_{\mathcal{P}}}$ possible replies to each sampled first message. For $k > 1$, Lemma A.1 and the hybrid argument replace this enumeration, which is why we only obtain a non-uniform circuit bound.

B Comparison with Existing Constructions

SNARGs for NP. Our lower bounds are consistent with all known constructions of SNARGs for NP from falsifiable assumption, each of which avoids them for a concrete reason.

- *Non-adaptive SNARGs from* iO *or* WE. Sahai and Waters [SW14] and Jin, Kalai, Lombardi, and Mathialagan [JKLM25] construct SNARGs for NP from iO for circuits or from witness encryption for NP. These assumptions are not efficiently falsifiable: checking that the adversary's challenge is valid requires deciding circuit equivalence or non-membership in L. Our results do not apply to such reductions, and in fact give a barrier to basing iO and WE themselves on efficiently falsifiable assumptions (Corollary 1.2).

- *Adaptive SNARGs from* iO. Adaptively sound SNARGs for NP [WW24, WW25] are likewise based on iO, and are not covered by our results for the same reason.

- *SNARGs from non-signaling PCPs.* Devadas, Hopkins, Kalai, Kothari, Lombardi, and Mathialagan [DHK+26] give a candidate construction based on non-signaling PCPs and a conjecture about Nullstellensatz proofs, which has since been refuted [Jin26]. Its security reduction uses advice that depends on the adversary. We show that such advice can be replaced by adversary-independent advice of size $2^{O(\ell + |\mathsf{crs}|)}$ (Section 7.3), which places L in $\mathsf{coAM}[2^{O(\ell+|\mathsf{crs}|)}\,\mathsf{poly}(\lambda)]$ (Corollary 7.10). Since their CRS is long, our results do not rule out their approach.

- *Reductions that are non-black-box in the adversary.* Jin, Lu, and Peng [JLP26] construct SNARKs for a class of UP languages from the polynomial hardness of LWE, with a security reduction that depends on the adversary's circuit. Hsieh, Jain, Li, and Mathialagan [HJLM26] construct SNARGs for NP from an assumption about the unprovability of certain mathematical statements, together with standard cryptographic and derandomization assumptions. These works use the SNARG adversary in a non-black-box way in their security arguments, so our lower bounds, which concern black-box reductions, do not apply.

Subclasses of NP. Our results also do not contradict SNARGs for subclasses of NP from efficiently falsifiable assumptions such as LWE, including SNARGs for P [CJJ22], for NTISP [BKK+18, KVZ21], for monotone policy batch NP [BBK+23], and for languages whose false statements have short propositional proofs of non-membership [JKLV24]. In each case, the complement of the language can already be decided or certified within time $2^{O(\ell)}\,\mathsf{poly}(n)$, so our conclusion holds for these languages unconditionally. For instance, $\mathsf{P} \subseteq \mathsf{coAM} \cap \mathsf{SIZE}[\mathsf{poly}]$. For NTISP, the proof is at

least as long as the space bound S, and the language is decidable in time $\mathsf{poly}(n) \cdot 2^S$ [BKK+18]. For monotone policy batch NP with k statements, each statement can be decided by brute force over its witnesses, and the policy then evaluated on the results, in total time $k \cdot 2^m \, \mathsf{poly}(n)$, where m is the length of a single witness; since the proof is at least m bits long, this is within $2^{O(\ell)} \, \mathsf{poly}(n)$. For [JKLV24], the reduction receives a short proof of non-membership as advice, so the language lies in coNP; in fact, this advice is exactly what the prover sends first in our protocol for the complement (Section 2.2).

Batch arguments and the tradeoff between CRS and proof length. Batch arguments (BARGs) illustrate an interesting phenomenon. A BARG proves that k statements $x_1, \ldots, x_k$ are all in L, with proofs that are short relative to k. From iO, there are fully succinct BARGs [GSWW22], whose proofs have length $\mathsf{poly}(\lambda)$ while the CRS grows with the length m of a single witness. From LWE, BARGs instead have a short CRS but proofs at least as long as a single witness; most recently, Döttling, Malavolta, and Paneth [DMP26] achieve a CRS independent of m and proofs of length $m + \mathsf{poly}(\lambda)$. Our lower bound suggests why the iO tradeoff has not been obtained from LWE. For $k = 1$, a BARG is a laconic argument for L, so by Theorem 1.1, a T-time one-query black-box reduction to a T-time falsifiable assumption places L in $\mathsf{coAM}[2^\ell \cdot \mathsf{poly}(T)]$, regardless of the CRS length. Under the hypothesis that SAT on m variables is not in $\mathsf{coAM}[2^{m^\delta}]$ for any constant $\delta < 1$, either $T \geq 2^{m^{1-o(1)}}$ or $\ell \geq m^{1-o(1)}$. In other words, unless the assumption is essentially 2^m-time falsifiable, the proof must be essentially as long as a single witness, however long the CRS is. In particular, the BARGs of [DMP26] are optimal up to the $o(1)$ in the exponent. In contrast, Chen and Jin [CJ25] only rule out the case where the CRS and the proof are both short. The iO-based BARGs avoid our lower bound by taking the other branch: iO is not efficiently falsifiable, and known instantiations from well-founded assumptions [JLS21] have $T \geq 2^m$.

Obfuscation. For obfuscation, known constructions of iO from well-founded assumptions [JLS21] rely on reductions running in time at least 2^m, where m is the input length of the obfuscated circuits. Thus, our bounds are trivial for these constructions, and hence also for the SNARGs of [SW14, JKLM25] once instantiated. The works [LZ17, KLMR18, JJ22, MDS25, JJMP25] that overcome the input-length barrier apply to restricted classes of circuit pairs, for instance those whose equivalence has short propositional proofs [JJ22, JJMP25]. For the circuits $C_{x,b}$ used in our lower bound, such a proof of equivalence is a proof that $x \notin L$, so our conclusion again holds trivially.