

On Soundness Notions for Interactive Oracle Proofs

¹, Alexander R. Block¹, Albert Garreta², Pratyush Ranjan Tiwari³, and Michał Zając²

¹Georgetown University and UMD College Park, alexander.r.block@gmail.com

²Nethermind, {albert,michal}@nethermind.io

³Johns Hopkins University, pratyush@cs.jhu.edu

Abstract

Interactive oracle proofs (IOPs) (Ben-Sasson et al., TCC 2016; Reingold et al., SICOMP 2021) have emerged as a powerful model for proof systems which generalizes both Interactive Proofs (IPs) and Probabilistically Checkable Proofs (PCPs). While IOPs are not any more powerful than PCPs from a complexity theory perspective, their potential to create succinct proofs and arguments has been demonstrated by many recent constructions achieving better parameters such as total proof length, alphabet size, and query complexity. In this work, we establish new results on the relationship between various notions of soundness for IOPs. First, we formally generalize the notion of round-by-round soundness (Canetti et al., STOC 2019) and round-by-round knowledge soundness (Chiesa et al., TCC 2019). Given this generalization, we then examine its relationship to the notions of generalized special soundness (Attema et al., CRYPTO 2021) and generalized special unsoundness (Attema et al., TCC 2022). We show that:

1. generalized special soundness implies generalized round-by-round soundness;
2. generalized round-by-round knowledge soundness implies generalized special soundness;
3. generalized special soundness does not imply generalized round-by-round knowledge soundness;
4. generalized round-by-round soundness (resp., special unsoundness) is an upper bound (resp., a lower bound) on standard soundness, and this relationship is tight when the round-by-round soundness and special unsoundness errors are equal; and
5. any special sound IOP can be transformed via (a variant of) the Fiat-Shamir transformation (in the Random Oracle Model) into a non-interactive proof that is adaptively sound in the Quantum Random Oracle Model.

1 Introduction

Probabilistic proof systems live at the heart of complexity theory and cryptography. Improvements in the practical efficiency of these proof systems have led to breakthroughs in zero-knowledge, delegation of computation, and other areas. Interactive oracle proofs (IOPs) were recently proposed [BCS16, RRR21] and have emerged as a powerful model for proof systems. Many recent constructions [BCS16, BBHR18, CMS19, GWC19, KPV19, COS20, BGKS20, Pol22, CBBZ23] of highly efficient and succinct proofs and arguments are compiled from IOPs. IOPs combine aspects of both probabilistically checkable proofs (PCPs) and interactive proofs (IPs), allowing a multi-round interaction between the prover and the verifier. A μ -round IOP can be viewed as a μ -round interactive proof (IP) where the verifier has PCP-like access to each prover message.

With the emergence of IOPs, the quest for understanding the security of IOPs has also started, specifically with respect to rendering IOPs non-interactive via the Fiat-Shamir transformation [FS87]. Fortunately, a recent fruitful line of work has introduced many tools to understand the security of IOPs under Fiat-Shamir: these include the notions of state-restoration soundness [BCS16], round-by-round soundness [CCH⁺19], and (generalized) special soundness [CDS94, Wik21, AFK22]. Another exciting line of work has attempted to establish relationships between these soundness notions for IPs [Hol19] and studied similar notions for IOPs [CMS19, KPV19, COS20].

This work formally analyzes and establishes the relationship between various soundness notions for IOPs. The first such notion of soundness is *round-by-round (RBR) soundness*, which captures the idea of “persistent falsehood” in an IOP: if the protocol is initiated in a state with a false statement and should be rejected by the verifier (i.e., a “doomed state”), then no matter how cleverly the prover responds in subsequent rounds, the protocol will “forever remain doomed” (except with negligible probability). *RBR knowledge soundness* captures the idea that there exists an (efficient) extractor algorithm such that if there was such a prover that could escape the “doomed state” with a higher (e.g., non-negligible) probability, then this extractor can extract a valid witness given the (partial) transcript of this interaction. Other key soundness notions we consider are *special soundness* and *special unsoundness* of IOPs. Special soundness was originally introduced in the context of Σ -protocols [Bab85, GMR89] and was later generalized in [CDS94, Wik18, Wik21]. A protocol is considered special sound if there exists an (efficient) extractor such that when given a tree of accepting transcripts (see Definition 3.5) as input, then the extractor can output a valid witness for the input. Special unsoundness [AFK22], on the other hand, argues about certain verifier challenges being extremely “lucky” for a malicious prover in the following sense: if the verifier sends such a challenge, then a malicious prover can convince the verifier of a false statement.

Establishing relationships among the various soundness notions for IOPs plays a critical role in proving the Fiat-Shamir security [FS87] of multi-round (i.e., non-constant round) protocols. Consequently, understanding these interrelations offers multiple avenues to ensure Fiat-Shamir security, significantly enhancing such cryptographic protocols’ applicability. This is demonstrated by our results (Figure 1), establishing new relationships among these soundness notions.

1.1 Our Results

In this paper, we formally establish new relations among generalized round-by-round (RBR) soundness and generalized special soundness. More formally, let \mathbf{R} be a relation (e.g., an NP relation) for which a μ -round interactive protocol is executed and let $L_{\mathbf{R}}$ be the language corresponding to the relation \mathbf{R} . Then for any statement $x \notin L_{\mathbf{R}}$, generalized RBR soundness is defined with respect a series of “doomed” sets \mathcal{D}_i for all $i = 0, \dots, \mu$. These sets represent protocol states (comprising of the statement and the transcript so far) from which the prover *cannot* possibly convince the verifier that $x \in L_{\mathbf{R}}$, except with small probability. Intuitively, the “doomed” set indicates that the protocol is in a point of no return for the prover: no matter what the prover does, except with small probability, the verifier will reject at the end of the interaction. Slightly more formally, generalized RBR soundness (Definition 3.11) requires the following:

- If a statement $x \notin L_{\mathbf{R}}$, then the protocol begins in this “doomed” state.
- If the current state (comprising the statement and the transcript of all messages so far) is “doomed”, then no matter what the prover’s next message is, the probability that the next state (including the prover’s next message and the verifier’s next message) is not doomed is at most $\varepsilon_i(|x|)$, where ε_i are predefined error functions. This means that once a state is doomed, it is highly likely that all future states will remain doomed, no matter what the prover does.

- After all the interaction rounds, if the interaction ends in a doomed state, then the verifier will reject.

Note that this generalization considers the errors in each round individually, unlike previous definitions as noted in [Definition 3.11](#).¹ Generalized RBR *knowledge soundness* is defined with respect to the same framework: a protocol is generalized RBR knowledge sound if there exists an efficient extractor such that if during any round i of the interaction, if the prover can escape the doomed set \mathcal{D}_i with probability larger than ε_i , then the extractor can extract a valid witness from the transcript of this interaction thus far. For generalized RBR knowledge soundness, the protocol *always* begins in a doomed state, even if $\mathfrak{x} \in L_{\mathbf{R}}$.

The generalized special soundness notion we consider is due to Attema et al. [[ACF21](#)] and is defined with respect to a tree of protocol transcripts. For a μ -round IOP $\Pi = (P, V)$, and $(k_1, \dots, k_\mu) \in \mathbb{N}^\mu$, a (k_1, \dots, k_μ) -*tree of accepting transcripts* for \mathfrak{x} is a set of $k = \prod_i k_i$ accepting transcripts $\{\tau_1, \dots, \tau_k\}$, each with common first message m , arranged in a tree of depth $\mu + 1$ as follows. The nodes in each tree correspond to the prover’s messages and the edges correspond to the verifier’s challenges, every node at depth $i - 1$ (for $1 \leq i \leq \mu$) has k_i children corresponding to pairwise distinct challenges, and every complete transcript corresponds to exactly one path from the root node to a leaf node in the tree. The protocol Π is (k_1, \dots, k_μ) -*special sound* if there exists a polynomial time algorithm Ext that, on input \mathfrak{x} and any (k_1, \dots, k_μ) -tree of accepting transcripts for \mathfrak{x} , outputs a witness \mathfrak{w} such that $(\mathfrak{x}, \mathfrak{w}) \in \mathbf{R}$.

The notion of ℓ -*special unsoundness* [[AFK22](#)] for an IOP $\Pi = (P, V)$ says that if there exists a dishonest prover strategy P^* such that during any round of the protocol, for any message m sent by P^* , there exists a “lucky” set of verifier challenges $\mathcal{L} \subset C$ of size $|\mathcal{L}| = \ell$ such that if the verifier V responds with $c \in \mathcal{L}$, then P^* can “behave honestly” for the remainder of the protocol and V will accept at the end of the protocol execution; here, C represents the set of verifier challenges.

For these generalized soundness notions, we prove the following results; see [Figure 1](#) for a high-level overview of all of our results. First we show that special soundness implies round-by-round soundness.

Theorem 1.1 (Special Soundness Implies RBR Soundness). *Let $\Pi = (P, V)$ be a μ -round IOP for a relation \mathbf{R} . Let C_i be the set of verifier challenges for round $i \in \{1, \dots, \mu\}$ and let $(k_1, \dots, k_\mu) \in \mathbb{N}^\mu$. Assume that Π is (k_1, \dots, k_μ) -special sound. Then Π is RBR sound with errors*

$$(\varepsilon_1, \dots, \varepsilon_\mu) = \left(\frac{k_1 - 1}{|C_1|}, \dots, \frac{k_\mu - 1}{|C_\mu|} \right). \quad (1)$$

Next, we show that round-by-round knowledge soundness implies special soundness.

Theorem 1.2 (RBR Knowledge Implies Special Soundness). *Let $\Pi = (P, V)$ be a μ -round IOP for a relation \mathbf{R} . Let C_i be the set of verifier challenges for round $i \in \{1, \dots, \mu\}$. Assume Π has round-by-round knowledge with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$, and let*

$$(k_1, \dots, k_\mu) = (\lceil |C_1| \varepsilon_1 \rceil + 1, \dots, \lceil |C_\mu| \varepsilon_\mu \rceil + 1).$$

Suppose $\sum_{i \in [\mu]} \prod_{j \in [i]} k_j$ is upper bounded by a polynomial (on the lengths of inputs). Then Π is (k_1, \dots, k_μ) -special sound.

We follow this up with a negative result: special soundness does not imply round-by-round knowledge soundness.

¹To the best of our knowledge, we are the first to formally define and analyze this generalized notion of round-by-round soundness. While this notion of RBR soundness has been implicit in prior works (e.g., in RBR soundness proofs of [[KPV19](#)] as one example), we were unable to find prior work formally defining and analyzing this generalized notion of RBR soundness.

Theorem 1.3 (Special Soundness does not Imply RBR Knowledge). *Assume $\text{NP} \neq \text{P}$. Then for any polynomial $\mu(|\mathbb{x}|)$, there exists a μ -round IOP Π with the following properties:*

- Π is $(1, \mu, 1)$ -special sound.
- If Π is RBR knowledge sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$, then $\varepsilon_i(\ell) = 1$ for some input length ℓ and some $i \in \{1, \dots, \mu\}$.

Finally, we show tight relationships between standard soundness, generalized round-by-round soundness, and special unsoundness.

Theorem 1.4 (Relation between Soundness, Round-by-round Soundness, and Special Unsoundness). *Let Π be a μ -round IOP for a relation \mathbf{R} with soundness error ε . Then the following hold:*

1. **RBR soundness is an upper bound for soundness.** *If Π is round-by-round sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$, then*

$$\varepsilon \leq 1 - \prod_{i=1}^{\mu} (1 - \varepsilon_i)$$

for all $\mathbb{x} \notin L_{\mathbf{R}}$.

2. **Special unsoundness is a lower bound for soundness.** *If Π is special unsound with errors $(\varepsilon'_1, \dots, \varepsilon'_\mu)$, then*

$$1 - \prod_{i=1}^{\mu} (1 - \varepsilon'_i) \leq \varepsilon$$

for all $\mathbb{x} \notin L_{\mathbf{R}}$. Moreover, there exists a dishonest unbounded prover P^* that, given any input \mathbb{x} , manages to make the verifier accept with probability at least $1 - \prod_i (1 - \varepsilon'_i)$.

3. **Tightness of RBR soundness, soundness, and special unsoundness.** *Suppose Π is round-by-round sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$ and that Π is special unsound with the same errors $(\varepsilon_1, \dots, \varepsilon_\mu)$. Then*

$$\varepsilon = 1 - \prod_{i=1}^{\mu} (1 - \varepsilon_i).$$

Moreover, the error is tight in the sense that there exists a dishonest prover P^* that, given any input \mathbb{x} , manages to have the verifier accept with probability at least ε .

1.1.1 Special Soundness and State-restoration Soundness

Our results on the relationship between special soundness and round-by-round soundness (i.e., [Theorem 1.1](#)) gives us new relationships between special soundness and *state-restoration soundness* [[BCS16](#)]. Informally, state-restoration (SR) soundness roughly states that an IOP remains secure (i.e., cannot convince a verifier of a false statement) even if a malicious prover is allowed to rewind the verifier to any prior state at most $b \geq 1$ times (see [[BCS16](#)] for complete details). It is known that state-restoration soundness and (non-generalized) round-by-round soundness are equivalent [[Hol19](#)] (up to some factors); moreover, state-restoration soundness error $\varepsilon_{\text{sr}}(b)$ and round-by-round soundness error ε_{rbr} must satisfy $\varepsilon_{\text{sr}}(b) \leq b\varepsilon_{\text{rbr}}$ [[CMS19](#), [KPV19](#), [COS20](#)], and this relation holds between state-restoration knowledge soundness and round-by-round knowledge soundness as well [[COS20](#)].

As a direct corollary of the above results (i.e., RBR soundness implies SR soundness) and [Theorem 1.1](#), we obtain the following result.

Corollary 1.5 (Special Soundness Implies State-restoration Soundness). *Let Π be a μ -round (k_1, \dots, k_μ) -special sound IOP with verifier challenge sets C_1, \dots, C_μ . Then for $b \geq 1$, Π has state-restoration soundness error*

$$\varepsilon_{\text{sr}}(b) \leq b \cdot \max_i \left\{ \frac{k_i - 1}{|C_i|} \right\}.$$

1.1.2 Special Soundness and Quantum-secure Fiat-Shamir

Recent cryptographic research put forth significant effort toward achieving post-quantum security of various cryptographic primitives, which include post-quantum security of non-interactive proofs obtained via the Fiat-Shamir transformation (or variants of this transformation) [CMS19, LZ19, DFMS19, CMSZ21]. With respect to post-quantum security, of interest to us is the so-called BCS transformation due to Ben-Sasson et al. [BCS16]; informally, this transformation compiles any IOP into a non-interactive proof via a variant of the Fiat-Shamir transformation in the random oracle model. In [BCS16], it is shown that applying this transformation to any state-restoration sound IOP results in an adaptively secure non-interactive proof in the random oracle model (ROM). The follow up work due to Chiesa et al. [CMS19] extend this result to show that compiling any round-by-round (knowledge) sound IOP with the BCS transformation yields an adaptively (knowledge) sound non-interactive proof in the ROM; we refer the reader to prior work (e.g., [BCS16, CMS19]) for complete details on this transformation. Furthermore, [CMS19] show that this transformation yields a non-interactive proof that is secure in the *quantum ROM* (QROM) (i.e., secure against quantum adversaries that are allowed to query the random oracle in superposition).

As a direct consequence of [CMS19] and [Theorem 1.1](#), any special sound IOP can be compiled via the BCS transformation to obtain an adaptively sound non-interactive proof in the QROM. This gives the following corollary.

Corollary 1.6 (Special Soundness Implies FS Security in the QROM). *Let Π be a μ -round (k_1, \dots, k_μ) -special sound IOP. Let $\text{BCS}(\Pi)$ be the non-interactive proof obtained by applying the BCS transformation to Π in the random oracle model, and let $\varepsilon = \max_i \{(k_i - 1)/|C_i|\}$. Then $\text{BCS}(\Pi)$ has adaptive soundness error $O(t^2\varepsilon + t^3/2^\lambda)$ against quantum attackers that make at most $t - O(q \log \ell)$ queries to the random oracle, where λ is the output length of the random oracle in bits, q is (an upper bound on) the total number of queries made by the verifier during any execution of Π , and ℓ is the total number of symbols sent by both the prover and verifier during any execution of Π .*

Remark 1.7. Note that [Corollary 1.6](#) directly implies that the Fiat-Shamir transformation of any special sound interactive proof in the ROM is an adaptively secure non-interactive proof in the QROM.

To the best of our knowledge, the above corollary is the first result relating the special soundness of a protocol and its security versus quantum adversaries when rendered non-interactive via the BCS transformation (i.e., a variant of the Fiat-Shamir transformation). Thus [Corollary 1.6](#) is the first result to our knowledge relating special soundness of multi-round protocols to quantum security. Notably, due to [Theorem 1.3](#), we *do not obtain* that the BCS transformation of a special sound IOP is knowledge sound versus quantum adversaries; we leave it as an interesting open problem to examine whether special soundness implies non-interactive knowledge soundness versus quantum adversaries. We remark that the Fiat-Shamir transformation of quantum-secure Σ -protocols (i.e., 1-round interactive arguments) was shown to be sound in the QROM [DFMS19].

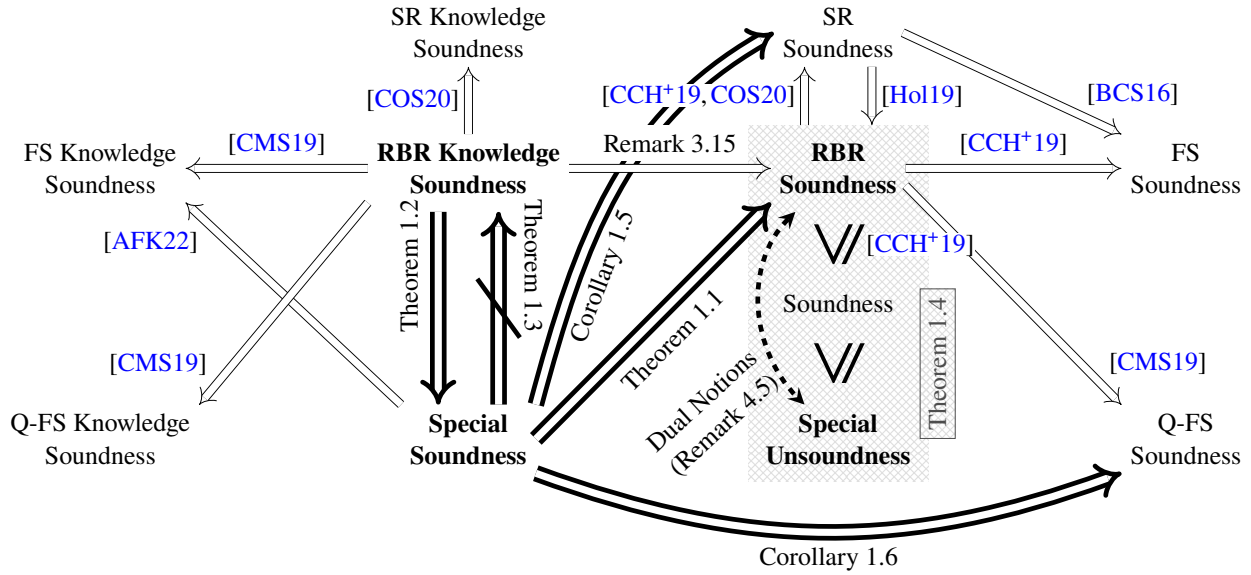


Figure 1: Pictorial overview of the relations between soundness notions. FS and Q-FS denote non-interactive adaptive security of the BCS transformation (i.e., a variant of the Fiat-Shamir transformation) in the random oracle model and quantum random oracle model, respectively. SR denotes state-restoration soundness. “ \Rightarrow ” arrows represent implications; “ $\not\Rightarrow$ ” represents there is no implication; and the dashed \leftrightarrow represents a relationship between notions (described by the text). Text in bold indicate the main soundness notions we study in this work. Thick arrows, lines, and background shading indicate our contributions. We remark that [Corollary 1.5](#) follows from [Theorem 1.1](#) and [\[CCH+19, COS20\]](#), and [Corollary 1.6](#) follows from [Theorem 1.1](#) and [\[CMS19\]](#).

1.2 Related Work

Interactive oracle proofs were introduced by Ben-Sasson et al. [BCS16] and independently by Reingold et al. [RRR21]. [BCS16] additionally introduced the notion of state-restoration soundness to formally show the Fiat-Shamir [FS87] security of multi-round IOPs in the random oracle model. The notion of round-by-round soundness was later introduced by Canetti et al. [CCH⁺19] for a similar reason: to prove Fiat-Shamir security of multi-round protocols, but in the *plain* model; however, round-by-round soundness readily implies Fiat-Shamir security in the random oracle model [CCH⁺19]. It was widely known that round-by-round soundness implies state-restoration soundness (e.g., [CCH⁺19, COS20]), and it was recently shown that state-restoration soundness implies round-by-round soundness [Hol19]. Special soundness was recently shown to also imply Fiat-Shamir security of multi-round protocols [AFK22]; moreover, this work also shows that special unsoundness of multi-round protocols readily admits an attack on the Fiat-Shamir transformed protocol.

Prior to these soundness tools, a variety of work [KRR17, CCRR18, HL18] circumvented the impossibility results of [BDG⁺13] by using stronger hardness assumptions to construct Fiat-Shamir compatible hash function families. Another line of work [GKR08, CMT12, BCGT13, Tha13, BTWV14, WTs⁺18, Set20, RR20] follows the frameworks of Kilian [Kil92] and Micali [Mic94] to compile interactive oracle proofs into efficient arguments and SNARKs via collision-resistant hash functions [Kil92, BCS16] or in the random oracle model [Mic94, BCS16].

2 Technical Overview

We give an overview of our main contributions in this section. Before we begin, we informally fix some notation. Given a μ -round IOP $\Pi = (P, V)$ for a relation \mathbf{R} and vector $(k_1, \dots, k_\mu) \in \mathbb{N}^\mu$, we say that tree T is a (k_1, \dots, k_μ) -tree of transcripts for Π on any input $\mathfrak{x} \in L_{\mathbf{R}}$ if T is a depth $\mu + 1$ tree such that nodes at level i have k_i outgoing edges, and the tree is labeled in the following way: every node at level i is labeled with a prover message and every edge is labeled with the corresponding challenge; moreover, the root of the tree is a single message m sent by the prover, and all leaves are prover messages. We say that T is an accepting tree of transcripts if all root to leaf paths are accepting transcripts (i.e., they are accepted by V). Now given the notion of a tree of accepting transcripts, the protocol Π is (k_1, \dots, k_μ) -special sound [ACK21] if there exists a polynomial time extractor algorithm that when given as input any (k_1, \dots, k_μ) -tree of accepting transcripts for an instance \mathfrak{x} , the extractor outputs a witness \mathfrak{w} such that $(\mathfrak{x}, \mathfrak{w}) \in \mathbf{R}$. For generalized round-by-round (knowledge) soundness, we point the reader back to the informal definition presented in Section 1.1 or the formal definition in Definitions 3.11 and 3.12.

2.1 Special Soundness Implies Round-by-round Soundness

In this section, we give an overview of Theorem 1.1: special soundness implies round-by-round soundness. More formally, we show that any (k_1, \dots, k_μ) -special sound IOP is generalized RBR sound with errors $\varepsilon_i = (k_i - 1)/|C_i|$. As a warm up, we first consider the non-generalized version of special soundness (i.e., k -ary trees of accepting transcripts) and RBR soundness (i.e., a single error bound ε and doomed set \mathcal{D} for all rounds of RBR soundness).

2.1.1 Warm Up: 1-round IOPs

Consider any 1-round (i.e., 3-move) public coin IOP $\Pi = (\mathsf{P}, \mathsf{V})$ with message space \mathcal{M} and challenge space C for some relation \mathbf{R} . Suppose that Π is k -special sound. That is, there is an extractor Ext such that on any input \mathbb{x} and any k -tree of accepting transcripts of the form (m, c_i, z_i) for $c_i \in C$ distinct for all i and $m, z_1, \dots, z_k \in \mathcal{M}$, the extractor outputs a witness w such that $(\mathbb{x}, w) \in \mathbf{R}$.

Now we claim that Π is round-by-round sound with error $\varepsilon = (k - 1)/|C|$. This can be seen as follows. Fix an input \mathbb{x} and consider any first message $m \in \mathcal{M}$ sent by the prover and any challenge c sent by the verifier. We say that (\mathbb{x}, m, c) is *completable* if there exists $z \in \mathcal{M}$ such that $\mathsf{V}(\mathbb{x}, m, c, z) = 1$; i.e., the verifier accepts. Intuitively speaking, for Π to be round-by-round sound, then the number of completable transcripts should be small for any $\mathbb{x} \notin L_{\mathbf{R}}$. In other words, for any first message m sent by the prover, the probability that (\mathbb{x}, m, c) is completable for $c \xleftarrow{\$} C$ should be at most ε defined above. Let $p(\mathbb{x}, m)$ be the probability that (\mathbb{x}, m, c) is completable for $c \xleftarrow{\$} C$, and suppose towards contradiction that $p(\mathbb{x}, m) > \varepsilon$. This implies that there exist $\alpha = \lceil |C| \cdot p(\mathbb{x}, m) \rceil > \varepsilon |C| = (k - 1)$ distinct challenges $c_1, \dots, c_\alpha \in C$ such that (\mathbb{x}, m, c_i) is completable for all $i \in \{1, \dots, \alpha\}$. Notice that $\alpha > k - 1$ and so $\alpha \geq k$. This tells us that any k of the α completable transcripts (\mathbb{x}, m, c_i) form a k -tree of accepting transcripts of the form $T = \{(\mathbb{x}, m, c_i, z_i)\}_i$ where z_i completes (\mathbb{x}, m, c_i) . Thus by k -special soundness of Π , we have that $(\mathbb{x}, \text{Ext}(\mathbb{x}, T)) \in \mathbf{R}$; however, this is a contradiction since $\mathbb{x} \notin L_{\mathbf{R}}$. Then it must be the case that at most $\alpha \leq (k - 1) = \varepsilon |C|$ distinct challenges can result in a completable transcript, which implies that $p(\mathbb{x}, m) \leq (k - 1)/|C|$ as desired.

More formally, to show that Π has round-by-round soundness error $\varepsilon = (k - 1)/|C|$, we define a suitable doomed set \mathcal{D} as follows. First consider the following two sets:

- Define \mathcal{D}_0 to be the set of all (\mathbb{x}, \emptyset) such that $\mathbb{x} \notin L_{\mathbf{R}}$, where \emptyset denotes the empty transcript.
- Define \mathcal{D}_1 to be the set of all (\mathbb{x}, m, c) for $m \in \mathcal{M}$ and $c \in C$ that are *not completable*.

Then we set $\mathcal{D} = \mathcal{D}_0 \cup \mathcal{D}_1$. Under this definition of \mathcal{D} , clearly we have $(\mathbb{x}, \emptyset) \in \mathcal{D}$ for all $\mathbb{x} \notin L_{\mathbf{R}}$ by definition; moreover, for any transcript (\mathbb{x}, m, c) that is not completable, then for all $z \in \mathcal{M}$ we have $\mathsf{V}(\mathbb{x}, m, c, z) = 0$. Now supposing that (\mathbb{x}, \emptyset) is doomed (by definition), consider any potential prover message m . Let $p(\mathbb{x}, m)$ denote the probability that $(\mathbb{x}, m, c) \notin \mathcal{D}$, where the probability is taken over $c \xleftarrow{\$} C$. Suppose there exists a message $m \in \mathcal{M}$ such that $p(\mathbb{x}, m) > \varepsilon$. Then by our above argument, we reach a contradiction as we can construct a k -tree of accepting transcripts and output a witness w , violating the assumption that $\mathbb{x} \notin L_{\mathbf{R}}$. Thus it must hold that $p(\mathbb{x}, m) \leq \varepsilon = (k - 1)/|C|$, as desired.

2.1.2 Extending to μ -round IOPs

Notice that the above argument crucially relies on the fact that if a malicious prover can leave the doomed set in the first round with probability larger than $\varepsilon = (k - 1)/|C|$, then we can manifest a k -tree of accepting transcripts. To extend the above argument to $\mu > 1$ round protocols, we want to preserve this fact. Recall the notion of a completable transcript from above, which states that for a given partial transcript (\mathbb{x}, m, c) , there exists a prover message $z \in \mathcal{M}$ that causes the verifier to accept the entire transcript (\mathbb{x}, m, c, z) . Now we extend this notion to any μ -round IOP. Consider any partial transcript of the form $\tau_i := (\mathbb{x}, m_1, c_1, \dots, m_i, c_i)$, where $i \leq \mu$ and message $m_j \in \mathcal{M}$ is sent by the prover at the start of round j and the verifier responds with challenge $c_j \xleftarrow{\$} C$, for all $j \leq i$. Now intuitively, we want to say that τ_i is completable if there is a sequence of messages and challenges $\sigma_{i+1} = (m_{i+1}, c_{i+1}, \dots, m_\mu, c_\mu, m_{\mu+1})$ such that $\mathsf{V}(\tau_i, \sigma_{i+1}) = 1$; i.e., (τ_i, σ_{i+1}) is a complete accepting transcript. However, this definition falls short of what we need to show our result.

In our argument for the 1-round IOP, we crucially relied on fact that the completable transcripts (\mathbb{x}, m, c) form an *entire* k -tree of accepting transcripts. Under our proposed extended definition of a completable transcript, this fact would no longer hold. To see this, suppose that $(\mathbb{x}, m_1, c_{1,i})$ is completable for at least k challenges $c_{1,1}, \dots, c_{1,k}$. Under our current definition of completable, this only implies that there exist k sequences of messages and challenges $\sigma_{1,i}$ such that $(\mathbb{x}, m_1, c_{1,i}, \sigma_{1,i})$ is an accepting transcript. Clearly, this is *not* a k -tree of accepting transcripts, so we can no longer derive our contradiction as with the 1-round IOP case.

We address the above issue by extending the definition of a completable transcript to a k -completable transcript. For any partial transcript τ_i for as defined above, we say that τ_i is k -completable if there exists a k -tree of transcripts T of depth $\mu + 1 - i$ such that for every $\sigma \in T$, the transcript (τ_i, σ) is accepted by the verifier, where σ is of the form $(m_{i+1}, c_{i+1}, \dots, m_\mu, c_\mu, m_{\mu+1})$. Now under this definition, we can now proceed to extend our proof to μ -round IOPs.

Suppose that Π is a k -special sound IOP. We argue that Π has round-by-round soundness error $\varepsilon = (k - 1)/|C|$. Define a doomed set \mathcal{D} as follows. First, consider the following sets:

- Define \mathcal{D}_0 to be the set of all (\mathbb{x}, \emptyset) such that $\mathbb{x} \notin L_{\mathbf{R}}$.
- For all $i \in \{1, \dots, \mu\}$, define \mathcal{D}_i to be the set of partial transcripts $\tau_i = (\mathbb{x}, m_1, c_1, \dots, m_i, c_i)$ such that τ_i is *not* k -completable.

Now take $\mathcal{D} = \bigcup_{i=0}^{\mu} \mathcal{D}_i$. Clearly, we have $(\mathbb{x}, \emptyset) \in \mathcal{D}$ for all $\mathbb{x} \notin L_{\mathbf{R}}$. Moreover, if $\tau_\mu = (\mathbb{x}, m_1, c_1, \dots, m_\mu, c_\mu)$ is not k -completable, then for all $m_{\mu+1} \in \mathcal{M}$ we have $\mathbb{V}(\tau_\mu, m_{\mu+1}) = 0$ by definition of k -completable.

Let $\tau_0 = (\mathbb{x}, \emptyset)$. We now argue that the probability that $\tau_1 = (\tau_0, m_1, c_1)$ is k -completable is at most ε for all $m_1 \in \mathcal{M}$ and $c_1 \xleftarrow{\$} C$. Our argument now proceeds identically to the 1-round IOP case. Let $p_1(\mathbb{x}, m_1)$ denote the probability over $c_1 \xleftarrow{\$} C$ that $\tau_1 = (m_1, c_1)$ is k -completable; in particular, $p_1(\mathbb{x}, m_1) = \Pr[(\mathbb{x}, \tau_1) \notin \mathcal{D}]$ by our definition. Suppose there exists $m_1 \in \mathcal{M}$ such that $p_1(\mathbb{x}, m_1) > \varepsilon$. Now this implies that there exist at least $\alpha > \varepsilon|C| = (k - 1)$ distinct challenges $c_{1,1}, \dots, c_{1,\alpha} \in C$ such that $(\mathbb{x}, m_1, c_{1,j})$ is k -completable. By definition of completable, this implies that for every $j \in \{1, \dots, \alpha\}$ and partial transcript $\tau_{1,j} = (m_1, c_{1,j})$, there exists a k -tree of transcripts $T_{2,j}$ of depth μ such that for all $\sigma \in T_{1,j}$, the complete transcript $(\mathbb{x}, \tau_{1,j}, \sigma)$ is accepted by the verifier. Since $\alpha \geq k$, taking any k out of α of the partial transcripts $(\mathbb{x}, m_1, c_{1,j})$ along with their k -tree of transcripts $T_{2,j}$ forms a k -tree of accepting transcripts of depth $\mu + 1$. Now by special soundness of Π , there exists an efficient extractor that extracts a witness \mathbb{w} such that $(\mathbb{x}, \mathbb{w}) \in \mathbf{R}$ when given this k -tree of accepting transcripts, contradicting our assumption that $\mathbb{x} \notin L_{\mathbf{R}}$. Thus it must be the case that $p_1(\mathbb{x}, m_1) \leq \varepsilon$, as required.

Now consider any intermediate round $i \leq \mu - 1$ and let $\tau_i = (m_1, c_1, \dots, m_i, c_i)$ be a partial transcript for this round. Suppose that $\tau_i \in \mathcal{D}$; we now show that for all $m_{i+1} \in \mathcal{M}$, the probability that $\tau_{i+1} = (\tau_i, m_{i+1}, c_{i+1})$ is k -completable is at most ε for $c_{i+1} \xleftarrow{\$} C$. Let $p_{i+1}(\mathbb{x}, \tau_i, m_{i+1}) = \Pr_{c_{i+1}}[(\mathbb{x}, \tau_i, m_{i+1}, c_{i+1}) \notin \mathcal{D}]$ denote this probability and suppose there exists $m_{i+1} \in \mathcal{M}$ such that $p_{i+1}(\mathbb{x}, \tau_i, m_{i+1}) > \varepsilon$. By the same argument as above and by definition of ε , there exist at least k distinct challenges $c_{i+1,1}, \dots, c_{i+1,k}$ such that the partial transcript $\tau_{i+1,j} = (\tau_i, m_{i+1}, c_{i+1,j})$ is k -completable. By definition of k -completable, this implies for each $\tau_{i+1,j}$ there exists a k -tree of transcripts $T_{i+2,j}$ of depth $\mu - i$ such that for all $\sigma \in T_{i+1,j}$, the transcript $(\tau_{i+1,j}, \sigma)$ is a complete and accepting transcript. Then we can construct a k -tree of transcripts T_{i+1} of depth $\mu - i + 1$ as $T_{i+1} = \{(\tau_i, m_{i+1}, c_{i+1,j}, \sigma)_{j \in \{1, \dots, k\}} : \sigma \in T_{i+1,j}\}$. Then clearly T_{i+1} forms a k -tree of transcripts that completes the partial transcript τ_i ; this contradicts the assumption that $\tau_i \in \mathcal{D}$, i.e., τ_i is not k -completable. Thus it must hold that $p_{i+1}(\mathbb{x}, \tau_i, m_{i+1}) \leq \varepsilon$ for all $m_{i+1} \in \mathcal{M}$, establishing the result.

2.1.3 Extending to Generalized Special Soundness and Generalized Round-by-round Soundness

The above argument naturally generalizes to (k_1, \dots, k_μ) -special soundness and $(\varepsilon_1, \dots, \varepsilon_\mu)$ round-by-round soundness with message spaces $\mathcal{M}_1, \dots, \mathcal{M}_{\mu+1}$ and challenge spaces $\mathcal{C}_1, \dots, \mathcal{C}_\mu$. First, we define a partial transcript τ_i to be (k_{i+1}, \dots, k_μ) -completable if there exists a (k_{i+1}, \dots, k_μ) -tree of transcripts T of depth $\mu - i + 1$ such that for all $\sigma \in T$, (τ_i, σ) is a complete transcript and $V(\mathbb{x}, \tau_i, \sigma) = 1$. Then under this definition, we define our doomed set $\mathcal{D} = \cup_{i=0}^\mu \mathcal{D}_i$ where \mathcal{D}_0 is defined identically as above and \mathcal{D}_i is the set of all partial transcripts τ_i that are not (k_{i+1}, \dots, k_μ) -completable. Then setting $\varepsilon_i = (k_i - 1)/|\mathcal{C}_i|$, we can proceed with the above argument in an identical fashion, deriving our contradictions and showing generalized round-by-round soundness of the IOP Π .

2.1.4 Special Soundness Implies Soundness in the Quantum Random Oracle Model

A direct consequence of the previous result is that the BCS transformation [BCS16] of a special sound IOP is sound in the Quantum Random Oracle Model (QROM). By our previous result, any special sound IOP is RBR sound, and the BCS transformation of any RBR sound IOP is a sound non-interactive proof in the QROM [CMS19], giving this consequence (i.e., Corollary 1.6). In particular, this implies that the Fiat-Shamir transformation of any special sound Interactive Proof is sound in the QROM.

2.2 Round-by-round Knowledge Implies Special Soundness

In this section, we give an overview of Theorem 1.2: round-by-round knowledge soundness implies special soundness.

2.2.1 Warm Up: 1-round IOPs

We again consider analyzing a 1-round public coin IOP $\Pi = (P, V)$ with message space \mathcal{M} and challenge space \mathcal{C} for some relation \mathbf{R} . Let ε be the round-by-round (RBR) knowledge error of Π . Let \mathbb{x} be arbitrary and let \mathcal{D} be the doomed set for RBR knowledge error. By RBR knowledge, there exists a polynomial time extractor Ext such that that if for all $m \in \mathcal{M}$ it holds that $p(\mathbb{x}, m) = \Pr_c[(\mathbb{x}, m, c) \notin \mathcal{D}] > \varepsilon$, then $\text{Ext}(\mathbb{x}, m)$ outputs a witness w such that $(\mathbb{x}, w) \in \mathbf{R}$, where the probability is taken over $c \xleftarrow{\$} \mathcal{C}$. Crucially, for all such transcripts $(\mathbb{x}, m, c) \notin \mathcal{D}$, there exists $z \in \mathcal{M}$ such that (\mathbb{x}, m, c, z) is an accepting transcript; we leverage this fact below.

Setting $k = \lceil |\mathcal{C}| \varepsilon \rceil + 1$, we now construct a polynomial-time extractor Ext' such that given any k -tree of accepting transcripts $T = \{(\mathbb{x}, m, c_i, z_i) : c_i \in \mathcal{C}, z_i \in \mathcal{M}\}_{i \in \{1, \dots, k\}}$, $(\mathbb{x}, \text{Ext}(T)) \in \mathbf{R}$. Suppose that for all $m \in \mathcal{M}$ we have $p(\mathbb{x}, m) > \varepsilon$. This implies that for any $m \in \mathcal{M}$, there are $\alpha > \varepsilon |\mathcal{C}|$ distinct challenges c_1, \dots, c_α such that $\text{Ext}(\mathbb{x}, m, c_i)$ outputs a valid witness, which implies that $\alpha \geq \lceil \varepsilon |\mathcal{C}| \rceil + 1 = k$, and thus there are at least k distinct challenges c_1, \dots, c_k such that $\text{Ext}(\mathbb{x}, m, c_i)$ outputs a valid witness for all $m \in \mathcal{M}$. Now given as input a tree of accepting transcripts T defined above, our new extractor $\text{Ext}'(\mathbb{x}, T)$ simply runs Ext on input (\mathbb{x}, m) , followed by $(\mathbb{x}, m, c_i, z_i)$ for all $i \in \{1, \dots, k\}$. Then Ext' outputs the same witness w that is given by Ext . Note that by assumption, we have that $(\mathbb{x}, m, c_i, z_i) \notin \mathcal{D}$, and we have $k = \lceil \varepsilon |\mathcal{C}| \rceil + 1$ such transcripts, so by RBR knowledge it must be the case that $\text{Ext}(\mathbb{x}, m, c_i, z_i)$ outputs a valid witness for some i . Otherwise, if $(\mathbb{x}, m, c_i, z_i) \in \mathcal{D}$, then the verifier would reject this transcript, contradicting T being a tree of accepting transcripts.

2.2.2 Extending to μ -round IOPs

Extending the above intuition to μ -round IOPs introduces some subtleties. Suppose again that Π is RBR knowledge sound with error ε and let $k = \lceil |\mathcal{C}|\varepsilon \rceil + 1$. Let T be a k -tree of accepting transcripts of depth $\mu + 1$ with root m_1 . We now construct a special soundness extractor that extracts a valid witness given the tree T .

Consider any partial transcript $\tau_i = (m_1, c_1, \dots, m_i, c_i)$ such that τ_i is a path in T . Then there is a unique message m_{i+1}^* such that (τ_i, m_{i+1}^*) is a path in T . Now suppose that τ_i has the following properties:

- $(\mathbb{x}, \tau_i) \in \mathcal{D}$ (i.e., τ_i is a doomed transcript);
- For all outgoing edges $c_{i+1,1}, \dots, c_{i+1,k}$ connected to m_{i+1}^* , we have $(\mathbb{x}, \tau_i, m_{i+1}^*, c_{i,j}) \notin \mathcal{D}$ for all j .

Clearly if τ_i has these properties, then our extractor Ext' on input T can run RBR knowledge extractor Ext on input $(\mathbb{x}, \tau_i, m_{i+1}^*, c_{i,j})$ and obtain a valid witness. This follows by our definition of k and by RBR knowledge: the fraction of challenges $c \in \mathcal{C}$ such that $(\mathbb{x}, \tau_i, m_{i+1}^*, c) \notin \mathcal{D}$ is strictly larger than ε ; i.e., $k/|\mathcal{C}| > \varepsilon$.

Now we claim that for *any* tree of accepting transcripts T , there exists some partial transcript τ_i that satisfies the above stated properties for some $i \in \{1, \dots, \mu\}$. If this claim is true, then the extractor Ext' on input any k -tree of accepting transcripts T simply runs Ext on input $(\mathbb{x}, \tau_j, m_{j+1}^*)$ over all possible partial transcripts τ_j that are in T that are a path from the root to node m_{j+1}^* . Then clearly Ext outputs a valid witness w such that $(\mathbb{x}, w) \in \mathcal{R}$; else this would violate RBR knowledge. Thus Ext' outputs this same witness w .

All that remains to be shown is the above claim, which we formally prove in [Section 4](#). At a high level, we show the claim by reverse (strong) induction on $i \in \{1, \dots, \mu\}$. For $i = \mu$, if the claim does not hold, then clearly T is no longer a tree of accepting transcripts as there is some complete transcript that remains in the doomed set, violating the assumption that T is an accepting tree of transcripts. Then fixing $i < \mu$ and assuming the claim is true for all j such that $i < j \leq \mu$, suppose τ_{i-1} is a partial transcript in the tree T . If m_i^* is the unique node in T such that $(\mathbb{x}, \tau_{i-1}, m_i^*)$ is a path in T and if for all outgoing edges $c_{i,1}, \dots, c_{i,k}$ of m_i^* we have $(\mathbb{x}, \tau_{i-1}, m_i^*, c_{i,j}) \notin \mathcal{D}$, then we are done. Otherwise, there exists some outgoing edge c_{i,j^*} of m_i^* such that $(\mathbb{x}, \tau_{i-1}, m_i^*, c_{i,j^*}) \in \mathcal{D}$. Now if this happens, by our induction hypothesis, we have that for m_{i+1}^* which has incoming edge c_{i,j^*} and all outgoing edges $c_{i+1,j}$ from m_{i+1}^* , the partial transcript $(\mathbb{x}, \tau_{i-1}, m_i^*, c_{i,j^*}, m_{i+1}^*, c_{i+1,j}) \notin \mathcal{D}$. This completes the induction step and the proof.

2.2.3 Extending to Generalized Round-by-round Knowledge and Generalized Special Soundness

The above argument again naturally extends to the generalized RBR knowledge and generalized special soundness cases. Taking $k_i = \lceil |\mathcal{C}_i|\varepsilon_i \rceil + 1$ for RBR knowledge errors ε_i , the above argument holds by replacing all k -tree of transcripts with (k_1, \dots, k_μ) -tree of transcripts and by considering the RBR knowledge error ε_i for any partial transcript τ_i . See [Section 4](#) for full details.

2.3 Special Soundness Does Not Imply Round-by-round Knowledge

Roughly stated, [Theorem 1.3](#) says that an IOP may be (k_1, \dots, k_μ) -special sound with small k_i 's, but at the same time only have RBR knowledge with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$ if $\varepsilon_i = 1$ for some i . We begin by providing intuition on why this is the case. To this end, observe that the special soundness extractor Ext_{spec} by definition is given access to a tree of accepting *complete* transcripts, while the RBR knowledge extractor Ext_{rbr} by definition only receives partial transcripts as inputs. Thus, in some sense Ext_{spec} has more information to work with than Ext_{rbr} . More precisely, suppose an IOP Π is built in a way that accepting complete transcripts “contain full information” about a valid witness, but that no partial transcript contains such information. Then, given a tree of accepting transcripts, the extractor Ext_{spec} is able to extract a witness from it, since the tree

contains complete transcripts that include “full information” about a valid witness. However, Ext_{rbr} is never given a complete accepting transcript, and instead only sees a partial transcript as its input. Consequently, by our assumptions on Π , the extractor Ext_{rbr} may be unable to reconstruct a witness from this partial transcript, no matter how likely it is that the partial transcript leaves certain doomed set. Under this informally described scenario, Π would be special sound, but it would not have RBR knowledge soundness.

We now formalize the above intuition. For simplicity, we restrict ourselves to IOP’s where the prover and verifier perform only one round of communication, i.e., $\mu(|\mathbb{x}|) = 1$ for all inputs \mathbb{x} . Our full result deals with IOP’s with arbitrarily (polynomially) many rounds of interaction (Section 4). Fix a relation $\mathbf{R} \in \mathbf{NP} \setminus \mathbf{P}$. Now we construct an IOP $\Pi = (P, V)$ for \mathbf{R} such that, as hinted above, its partial transcripts contain “no information” about witnesses, while its complete transcripts do. To this end, for any $(\mathbb{x}, \mathbb{w}) \in \mathbf{R}$, we have P send a single bit 0 as its first message m_1 . Then we have V reply with a random string c , to which the honest prover P replies by sending the entire witness \mathbb{w} . Then V accepts the proof if and only if $(\mathbb{x}, \mathbb{w}) \in \mathbf{R}$ and $m_1 = 0$. Clearly, Π is 1-special sound since any complete accepting transcript for \mathbb{x} contains a witness \mathbb{w} . On the other hand, all non-complete partial transcripts for Π are of the form (\mathbb{x}) , $(\mathbb{x}, 0)$, or $(\mathbb{x}, 0, c)$. None of these partial transcripts have information related to a witness for \mathbb{x} (other than the information provided by the input \mathbb{x} itself).

Now assume Π has RBR knowledge with error ε_1 , and let \mathcal{D} and Ext be a corresponding doomed set and an extractor algorithm. Then, by definition of RBR knowledge, whenever we have

$$\Pr_c[(\mathbb{x}, 0, c) \notin \mathcal{D}] > \varepsilon_1, \quad (2)$$

Ext can extract a valid witness for \mathbb{x} just from seeing $(\mathbb{x}, 0)$. However, if $\mathbb{x} \in L_{\mathbf{R}}$, then the probability $\Pr_c[(\mathbb{x}, 0, c) \notin \mathcal{D}]$ must be 1 since for all c , the partial transcript $(\mathbb{x}, 0, c)$, which comprises all μ rounds of interaction, can be extended into a complete accepting transcript $(\mathbb{x}, 0, c, m)$. Then by definition of doomed set, we must have $(\mathbb{x}, 0, c) \notin \mathcal{D}$. It follows from this and Eq. (2) that if $\varepsilon_1(|\mathbb{x}|) < 1$, then Ext is able to output a witness for \mathbb{x} just from seeing $(\mathbb{x}, 0)$. With these arguments in mind and assuming $\varepsilon_1(|\mathbb{x}|) < 1$, it is straightforward to build a deterministic polynomial time algorithm that recognizes the language $L_{\mathbf{R}}$, contradicting our assumption that \mathbf{R} is in \mathbf{NP} but not in \mathbf{P} . Thus, Π cannot have RBR knowledge soundness; i.e., if it has RBR knowledge, then the error ε_1 is not negligible (in fact, it is 1).

In Section 4, we formally present the above ideas. Moreover, we generalize these arguments to construct a μ -round IOP that is special sound but does not have RBR knowledge soundness for any polynomial $\mu = \mu(|\mathbb{x}|)$.

2.4 Special Unsoundness and Round-by-round Soundness are Dual

Finally, we relate round-by-round soundness and the notion of *special unsoundness* [AFK22]. Recall that, informally, an IOP $\Pi = (P, V)$ is ℓ -special unsound if there exists a dishonest prover strategy P^* such that during any round of the protocol, for any message m sent by P^* , there exists a “lucky” set of verifier challenges $\mathcal{L} \subset \mathcal{C}$ such that $|\mathcal{L}| = \ell$ such that if the verifier V responds with $c \in \mathcal{L}$, then P^* can “behave honestly” for the remainder of the protocol and V will accept at the end of the protocol execution.

Given the above notion of special unsoundness, it is immediate on an intuitive level that special unsoundness and RBR soundness are “dual” notions: RBR soundness states that for any dishonest prover strategy, the probability the prover gets “lucky” is upper bounded by some ε , whereas special unsoundness says that there exists a prover strategy where the probability the prover gets “lucky” is lower bounded by some ε' . We formalize this relationship in Section 4.1. Assume that Π is an ε -sound μ -round IOP, and assume that Π is generalized round-by-round sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$ and special unsound with errors $(\varepsilon'_1, \dots, \varepsilon'_\mu)$. Then for $\varepsilon = 1 - \prod_i (1 - \varepsilon_i)$ and $\varepsilon' = 1 - \prod_i (1 - \varepsilon'_i)$, it holds that $\varepsilon' \leq \varepsilon \leq \varepsilon$. Moreover, we show that if $\varepsilon'_i = \varepsilon_i$ for all i , then $\varepsilon = \varepsilon$. See Section 4.1 for complete details.

3 Preliminaries

A relation \mathbf{R} is a subset of pairs $(\mathbb{x}; \mathbb{w}) \in \{0, 1\}^* \times \{0, 1\}^*$. The strings \mathbb{x} are called *inputs* (these are often called also *statements* or *instances*), and the strings \mathbb{w} are called *witnesses*. To each relation \mathbf{R} there corresponds a language $L_{\mathbf{R}} \subseteq \{0, 1\}^*$ consisting of all statements \mathbb{x} such that $(\mathbb{x}, \mathbb{w}) \in \mathbf{R}$ for some \mathbb{w} ; i.e., $L_{\mathbf{R}} := \{\mathbb{x} : \exists \mathbb{w} \text{ s.t. } (\mathbb{x}, \mathbb{w}) \in \mathbf{R}\}$. When $(\mathbb{x}, \mathbb{w}) \in \mathbf{R}$, we say that \mathbb{w} is a *valid witness* for \mathbb{x} . We assume our relations to be in the class **NP**.

We parameterize our security functions either by the length $|\mathbb{x}|$ of an input, but for ease of exposition we omit these from our notation, i.e., we write expressions such as “soundness error ε ” instead of “soundness error $\varepsilon(|\mathbb{x}|)$ ”. We proceed similarly for other types of functions. A function ε is *negligible* if $\varepsilon(|\mathbb{x}|) = o(1/p(|\mathbb{x}|))$ for any positive polynomial p .

We denote by \mathbb{N} be the set of all positive integers. For any $m \in \mathbb{N}$, we let $[m]$ denote the set $\{1, \dots, m\}$. For any finite set S , we let $s \xleftarrow{\$} S$ denote the process of sampling an element of S uniformly and independently at random.

3.1 Interactive Oracle Proofs

Given a map $f \in A^B$ for some sets A, B , we denote by $\llbracket f \rrbracket$ an *oracle* to the map f . This is a hypothetical algorithm that takes elements $a \in A$ as input, and outputs $f(a)$ instantaneously.

Definition 3.1 (Interactive Proofs (IP)). *A μ -round interactive proof for a relation \mathbf{R} is a pair of interactive algorithms $\Pi = (P, V)$ such that:*

- For $\mathbb{x} \in L_{\mathbf{R}}$ and \mathbb{w} such that $(\mathbb{x}, \mathbb{w}) \in \mathbf{R}$, before the start of the protocol, P receives both (\mathbb{x}, \mathbb{w}) as input and V receives \mathbb{x} as input.
- $P(\mathbb{x}, \mathbb{w})$ and $V(\mathbb{x})$ exchange $2\mu(|\mathbb{x}|) + 1$ messages, where P sends the first and last message, and during any round of interaction P sends message m_i to V . After P sends $m_{\mu(|\mathbb{x}|)+1}$, V either accepts (outputs 1) or rejects (outputs 0).

We require the following properties to hold:

- **Completeness:** for all $(\mathbb{x}, \mathbb{w}) \in \mathbf{R}$, we have

$$\Pr [\langle P(\mathbb{w}), V \rangle(\mathbb{x}) = 1] = 1,$$

where $\langle P(\mathbb{w}), V \rangle(\mathbb{x})$ denotes the output of P and V interacting on common input \mathbb{x} where P is additionally given \mathbb{w} as input, and the above probability holds over the random coins of V .

- **ϵ -Soundness:** for any $\mathbb{x} \notin L_{\mathbf{R}}$ and any unbounded interactive algorithm P^* , we have

$$\Pr [\langle P^*, V \rangle(\mathbb{x}) = 1] \leq \epsilon,$$

where the probability is taken over the random coins of V .

We say that Π is public-coin if all messages sent by V are independent uniform random strings of some bounded length and the output of V does not depend on any secret state.

Remark 3.2. In this paper, all IOP’s are assumed to be public-coin. Also, we often write μ to denote both the function μ and the value $\mu(|\mathbb{x}|)$.

Definition 3.3 (Interactive Oracle Proof). An *Interactive Oracle Proof (IOP)* for a relation \mathbf{R} is a μ -round IP (P, V) for \mathbf{R} such that for all \mathfrak{x} , at the start of each round of interaction $i \in [\mu(|\mathfrak{x}|)]$, P sends m_i and V receives oracle access to m_i via $\llbracket m_i \rrbracket$. Crucially, at the end of the interactive phase, V does not necessarily need to read the whole m_i in order to decide whether to accept or reject.

Definition 3.4 (Message Spaces and i -round Partial Transcripts). Let Π be a μ -round IOP $\Pi = (P, V)$ for a relation \mathbf{R} . We let $\mathcal{M}_1, \dots, \mathcal{M}_{\mu+1}$ denote the sets of all potential prover's messages such that in any round $i \in [\mu]$, P sends a message from \mathcal{M}_i . The set $\mathcal{M}_{\mu+1}$ is the set of all potential prover's final messages. Similarly, we let $\mathcal{C}_1, \dots, \mathcal{C}_\mu$ denote the sets of all potential verifier's challenges such that in any round $i \in [\mu]$, V sends a challenge from \mathcal{C}_i sampled uniformly at random.

For $i \in [\mu]$, we define a i -round partial transcript as a vector of the form $\tau = (m_1, c_1, \dots, m_i, c_i)$ where $m_i \in \mathcal{M}_j$ and $c_i \in \mathcal{C}_j$ for all $j \in [i]$. We also let a 0-round partial transcript be a vector of the form (\mathfrak{x}) . We write $\text{PartTr}(i)$ to denote the set of all i -round partial transcripts. A complete transcript is a transcript of the form (\mathfrak{x}, τ, m) where $(\mathfrak{x}, \tau) \in \text{PartTr}(\mu)$ and $m \in \mathcal{M}_{\mu+1}$. Such a transcript is accepting if $V(\mathfrak{x}, \tau, m) = 1$.

3.2 Special Soundness

Before introducing the notion of special soundness, we introduce notions related to *trees of transcripts*.

Definition 3.5 ((k_1, \dots, k_μ) -Tree of Transcripts). Let $\Pi = (P, V)$ be a μ -round IOP. Let $(k_1, \dots, k_\mu) \in \mathbb{N}^\mu$. A (k_1, \dots, k_μ) -tree of transcripts for \mathfrak{x} is a set of $k = \prod_i k_i$ complete transcripts (τ_1, \dots, τ_k) with common first message m , arranged in a tree of depth $\mu + 1$ ² and arity k_1, \dots, k_μ , respectively. The nodes in the tree correspond to the prover's messages, and the edges correspond to the verifier's challenges. Every internal node at depth $i - 1$ ($1 \leq i \leq \mu$) has k_i children with distinct challenges. For every $j \in [k]$, τ_j represents one root-to-leaf path. Finally, we say that the tree is a (k_1, \dots, k_μ) -tree of accepting transcripts for \mathfrak{x} if every transcript is accepted by V .

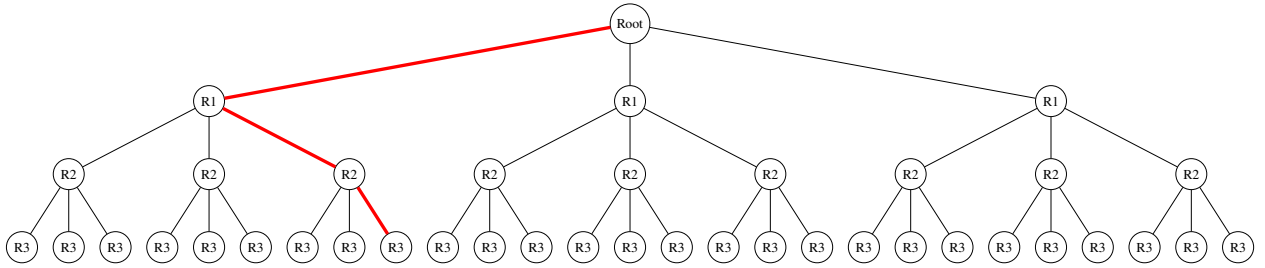


Figure 2: $(3, 3, 3)$ -tree of transcripts with a highlighted complete/accepting transcript.

Given Definition 3.5, we can now define special soundness.

Definition 3.6 (Special Soundness). Let Π be a μ -round IOP for a relation \mathbf{R} , and let $(k_1, \dots, k_\mu) \in \mathbb{N}^\mu$. We say Π is (k_1, \dots, k_μ) -special sound if there exists a polynomial time algorithm Ext that, on input \mathfrak{x} and any (k_1, \dots, k_μ) -tree of accepting transcripts for \mathfrak{x} , outputs a witness \mathfrak{w} such that $(\mathfrak{x}, \mathfrak{w}) \in \mathbf{R}$.

²We set the root vertex of a tree to be of depth 1 (as opposed to 0).

3.2.1 Notions Related to Trees of Transcripts

We need a few more technical definitions related to [Definition 3.5](#) for our proofs. One notion is that of a “partial” tree of accepting transcripts.

Definition 3.7 ((k_i, \dots, k_μ) -tree of Transcripts). *Let Π be a μ -round IOP, let $(k_1, \dots, k_\mu) \in \mathbb{N}^\mu$, and let $i \in [\mu]$. By a (k_i, \dots, k_μ) -tree of transcripts we refer to a subtree T' of a (k_1, \dots, k_μ) -tree T of transcripts such that the root node of T' is at depth i in T . In other words, T' is a tree of depth $\mu - i + 1$, where for depths $1, \dots, \mu - i + 1$, nodes have k_i, \dots, k_μ children, respectively. Each complete path in T' corresponds to a suffix of a complete transcript. Note that if $i = 1$ then we recover the original definition of (k_1, \dots, k_μ) -tree of transcripts.*

We also need the notions of *complete and rooted paths* along with a way to identify paths via their labeling.

Definition 3.8 (Complete Paths, Rooted Paths, and Identification of Paths via their Label). *Let T be a tree of accepting transcripts. We say a path in T is complete if it starts at the root node and it ends in a leaf node. We say a path is rooted if it start at the root node.*

For ease of presentation we introduce the following abuse of notation and terminology: we identify rooted paths via their labels, which are partial transcripts.³

Finally, we need the notion of a *completable transcript*.

Definition 3.9 (Completable Transcripts). *Let $i \in [\mu + 1]$ and let $(\mathbb{x}, \tau) \in \text{PartTr}(i - 1)$ (see [Definition 3.4](#)). We say (\mathbb{x}, τ) is completable if one of the following holds:*

- $i \leq \mu$ and there exists a (k_i, \dots, k_μ) -tree of transcripts T' such that, for all complete path τ' in T' , the transcript $(\mathbb{x}, \tau, \tau')$ is accepted by the verifier (following our convention from [Definition 3.8](#), here we identify the path τ' with the sequence of messages and challenges associated to each of its nodes and edges (i.e., its label)). We say that T' completes (\mathbb{x}, τ) .
- $i = \mu$ and there exists $m \in \mathcal{M}_{\mu+1}$ such that the complete transcript (\mathbb{x}, τ, m) is accepted by the verifier.

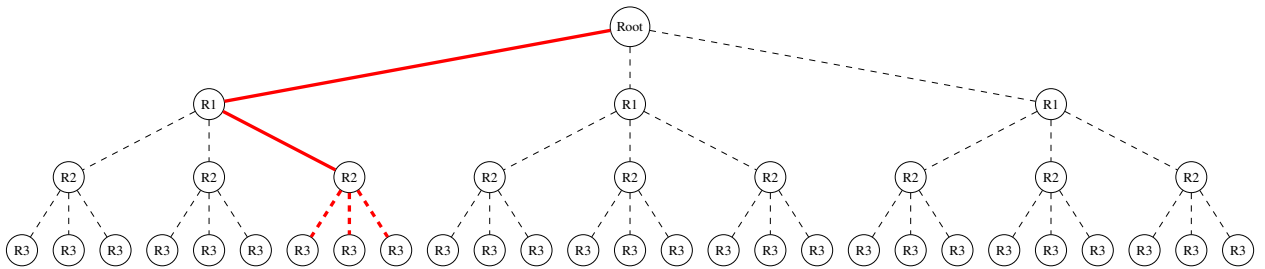


Figure 3: 3-round IOP tree with a highlighted completable partial transcript. In this figure, the highlighted path through the tree ends in the second round with node R2, and the hanging tree (with dashed red edges) at R2 represents the tree that completes the partial transcript.

³Technically, this is ambiguous because two distinct rooted paths could have the same label. However in our context there will never be risk of confusion.

3.3 Special Unsoundness

The notion of special unsoundness was introduced in [AFK22] in the context of analyzing the security of the Fiat-Shamir transformation of the parallel repetition of an interactive proof. In particular, the authors describe an attack on the Fiat-Shamir transformed protocol if it is special unsound.

Definition 3.10 (Special Unsoundness [AFK22]). *Let Π be a μ -round IOP, and let $(\ell_1, \dots, \ell_\mu) \in \mathbb{N}^\mu$. We say that Π has $(\ell_1, \dots, \ell_\mu)$ -special unsoundness if there exists a dishonest prover \mathcal{A} of the following form and, so that in the execution with V and input \mathbb{x} the following holds:*

- \mathcal{A} starts off in active mode, which is so that in every round i , when \mathcal{A} sends the message m_i , there exists a subset $\mathcal{L}_i \subseteq C_i$ such that $|\mathcal{L}_i| = \ell_i$ (defined as a function of the state of \mathcal{A} at that point) such that if the subsequent challenge c_i is in \mathcal{L}_i , then \mathcal{A} switches into passive mode.
- If \mathcal{A} switches into passive mode, then it remains in passive mode until the end of the protocol, and V accepts at the end of the protocol.

3.4 Generalized Round-by-round (Knowledge) Soundness

We define the notions of *generalized round-by-round soundness* and *generalized round-by-round knowledge soundness*, respectively. These are essentially the original definitions from [CCH⁺19, CMS19] with the following modification: instead of having a “one-size-fits-all” soundness/knowledge error for all rounds of the protocol, we consider individual errors for each round of the protocol. Following [Hol19], we use a formalism based on “doomed” sets of partial transcripts, as opposed to using a “state function” in the original formulation (see [CCH⁺19, Hol19] for details).

Definition 3.11 ((Generalized) Round-by-round Soundness). *An IOP $\Pi = (P, V)$ for a relation \mathbf{R} has (generalized) round-by-round soundness with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$ if there exists a (not necessarily efficiently computable) “doomed set” \mathcal{D} of partial transcripts such that the following properties hold:*

1. *If $\mathbb{x} \notin L_{\mathbf{R}}$, then $(\mathbb{x}, \emptyset) \in \mathcal{D}$, where \emptyset denotes the empty transcript.*
2. *For any \mathbb{x} and any μ -round partial transcript $(\mathbb{x}, \tau) \in \text{PartTr}(\mu)$ and any last prover message $m \in \mathcal{M}_{\mu+1}$, if $(\mathbb{x}, \tau) \in \mathcal{D}$ then $V(\mathbb{x}, \tau, m) = 0$.*
3. *If $(\mathbb{x}, \tau) \in \text{PartTr}(i-1)$ is a $(i-1)$ -round partial transcript for some $i \in [\mu]$ and $(\mathbb{x}, \tau) \in \mathcal{D}$, then for all $m \in \mathcal{M}_i$ we have*

$$\Pr_{\substack{c \leftarrow C_i \\ \tau, m}} [(\tau, m, c) \notin \mathcal{D}] \leq \varepsilon_i.$$

If for all $i \in [\mu]$, $\varepsilon_i = \varepsilon_i(|\mathbb{x}|)$ is a negligible function of the input length $|\mathbb{x}|$, then we simply say that Π has RBR soundness.

The original definition of round-by-round soundness considers only the scenario in which $\varepsilon_1 = \dots = \varepsilon_\mu$, and in that case one says that Π has *RBR soundness with error ε* , where $\varepsilon = \varepsilon_i$ for all i .

Definition 3.12 ((Generalized) Round-by-round Knowledge). *We say an IOP $\Pi = (P, V)$ for a relation \mathbf{R} has (generalized) round-by-round knowledge with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$ if there exists a (not necessarily efficiently computable) “doomed set” \mathcal{D} of partial transcripts such that the following properties hold:*

- For all possible inputs \mathbb{x} , we have $(\mathbb{x}, \emptyset) \in \mathcal{D}$.
- For any μ -round partial transcript $(\mathbb{x}, \tau) \in \text{PartTr}(\mu)$ and any last prover message $m \in \mathcal{M}_{\mu+1}$, if $(\mathbb{x}, \tau) \in \mathcal{D}$, then $\mathbb{V}(\mathbb{x}, \tau, m) = 0$.
- There exists a polynomial time algorithm Ext , called extractor, with the following properties. If $(\mathbb{x}, \tau) \in \text{PartTr}(i-1)$ for some $i \in [\mu]$, $(\mathbb{x}, \tau) \in \mathcal{D}$, and for all $m \in \mathcal{M}_i$ we have

$$\Pr_{c \stackrel{\$}{\leftarrow} \mathcal{C}_i} [(\mathbb{x}, \tau, m, c) \notin \mathcal{D}] > \varepsilon_i,$$

then $\text{Ext}(\mathbb{x}, \tau, m)$ outputs a witness w such that $(\mathbb{x}, w) \in \mathbf{R}$. In this case, we say that the partial transcript (\mathbb{x}, τ, m) above is a RBR extractable partial transcript.

If for all $i \in [\mu]$, $\varepsilon_i = \varepsilon_i(|\mathbb{x}|)$ is a negligible function of the input length $|\mathbb{x}|$, then we simply say that Π has RBR knowledge soundness.

Remark 3.13. We often drop “generalized” when discussing the above definitions and simply write RBR soundness and RBR knowledge, unless otherwise stated (e.g., for clarity).

Remark 3.14. Let Π be an IOP with μ rounds. Suppose Π is round-by-round sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$. Then Π is round-by-round sound with error ε , where

$$\varepsilon = \max_{i \in [\mu]} \{\varepsilon_i\}.$$

Conversely, if Π is round-by-round sound with error ε , then it is generalized round-by-round sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$ where $\varepsilon_i = \varepsilon$ for all $i \in [\mu]$. Note that the original definitions of RBR soundness and knowledge [CCH⁺19, CMS19] exactly consider this case (i.e., all errors are identical).

Remark 3.15 (RBR Knowledge Implies RBR Soundness). An IOP with RBR knowledge with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$ is necessarily RBR sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$. This is because if (\mathbb{x}, τ, m) is an RBR extractable partial transcript, then the extractor from the definition of RBR knowledge outputs a witness w such that $(\mathbb{x}, w) \in \mathbf{R}$, and so $\mathbb{x} \in L_{\mathbf{R}}$. Hence, if \mathcal{D} is the doomed set with which Π has RBR knowledge, the subset $\mathcal{D}' \subseteq \mathcal{D}$ consisting of all $(\mathbb{x}, \tau) \in \mathcal{D}$ such that $\mathbb{x} \notin L_{\mathbf{R}}$ is a doomed set with which Π has RBR soundness with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$.

4 Relating Special Soundness and Round-by-round Knowledge

In this section, we recall our main results one by one and prove them in turn. We first recall [Theorem 1.1](#).

Theorem 1.1 (Special Soundness Implies RBR Soundness). *Let $\Pi = (P, V)$ be a μ -round IOP for a relation \mathbf{R} . Let \mathcal{C}_i be the set of verifier challenges for round $i \in \{1, \dots, \mu\}$ and let $(k_1, \dots, k_\mu) \in \mathbb{N}^\mu$. Assume that Π is (k_1, \dots, k_μ) -special sound. Then Π is RBR sound with errors*

$$(\varepsilon_1, \dots, \varepsilon_\mu) = \left(\frac{k_1 - 1}{|\mathcal{C}_1|}, \dots, \frac{k_\mu - 1}{|\mathcal{C}_\mu|} \right). \quad (1)$$

Proof. Assume Π is (k_1, \dots, k_μ) -special sound. We define a doomed set \mathcal{D} of partial transcripts as the union $\mathcal{D} = \bigcup_{i=0}^\mu \mathcal{D}_i$, where each $\mathcal{D}_i \subseteq \text{PartTr}(i)$, i.e., \mathcal{D}_i consists of i -round partial transcripts. These sets \mathcal{D}_i are defined as follows.

$$\begin{aligned} \mathcal{D}_0 &:= \{(\mathbb{x}) \mid \mathbb{x} \notin L_R\}, \\ \mathcal{D}_i &:= \{(\mathbb{x}, \tau) \in \text{PartTr}(i) \mid (\mathbb{x}, \tau) \text{ is not completable}\}. \end{aligned}$$

We now prove that Π has RBR soundness with errors $\varepsilon_i = (k_i - 1)/|C_i|$ for all $i \in [\mu]$. Let Ext be the special soundness extractor and fix an input \mathbb{x} . Let $i \in [\mu]$ and let $(\mathbb{x}, \tau) \in \text{PartTr}(i-1) \cap \mathcal{D} = \mathcal{D}_{i-1}$ be a doomed $(i-1)$ -round partial transcript. In particular, either $i = 1$ and $(\mathbb{x}, \tau) = (\mathbb{x}) \in \mathcal{D}_0$, or $(\mathbb{x}, \tau) \in \mathcal{D}_i$ for $i > 1$, and (\mathbb{x}, τ) is not completable. Given $m \in \mathcal{M}_i$, let

$$p_i(\mathbb{x}, \tau, m) := \Pr_{c \stackrel{\$}{\leftarrow} C_i} [(\mathbb{x}, \tau, m, c) \notin \mathcal{D}_i]$$

denote the probability of leaving the doomed set \mathcal{D}_i . Assume $p_i(\mathbb{x}, \tau, m) > \varepsilon_i$. Then there are at least k_i distinct challenges c_1, \dots, c_{k_i} such that $(\mathbb{x}, \tau, m, c_j) \notin \mathcal{D}_i$ for all $j \in [k_i]$. Thus, it follows that each $(\mathbb{x}, \tau, m, c_j)$ is completable. We now claim that (\mathbb{x}, τ) is also completable. Let T_1, \dots, T_{k_i} each be a (k_{i+1}, \dots, k_μ) -tree of transcripts such that each tree T_j completes the i -round partial transcript $(\mathbb{x}, \tau, m, c_j)$ for all $j \in [k_i]$. By definition, for each complete path τ_j in T_j , the complete transcript $(\mathbb{x}, \tau, m, c_j, \tau_j)$ is accepted by the verifier. Now construct a (k_i, \dots, k_μ) -tree T of transcripts as follows: create a root node V_R with label m and create k_i children $V_{R_1}, \dots, V_{R_{k_i}}$ for the root node V_R , labeling the edges with the challenges c_1, \dots, c_{k_i} , respectively. Now attach the trees T_1, \dots, T_{k_i} to the child vertices $V_{R_1}, \dots, V_{R_{k_i}}$, respectively. The resulting tree T is a (k_i, \dots, k_μ) -tree of transcripts with root labeled as m . Moreover, each complete path in T has label of the form (m, c_j, τ') for some $j \in [k_i]$. Hence, T completes (\mathbb{x}, τ) since $(\mathbb{x}, \tau, m, c_j, \tau')$ is a complete transcript that is accepted by the verifier for all choices of m, c_j and τ' .

Given that (\mathbb{x}, τ) is completable, if $i \geq 2$ then $p_i(\mathbb{x}, \tau, m)$ cannot be larger than ε_i if $(\mathbb{x}, \tau) \in \mathcal{D}_{i-1}$. Now assume that $i = 1$ so that $(\mathbb{x}, \tau) = (\mathbb{x}, \emptyset) \in \text{PartTr}(0)$ is a 0-round partial transcript. By definition, we have $(\mathbb{x}, \emptyset) \in \mathcal{D}_0$. Suppose towards contradiction that $p_1(\mathbb{x}, m) > \varepsilon_1$ for some $m \in \mathcal{M}_1$. Then the tree T with root m that completes (\mathbb{x}, \emptyset) (obtained through the argument above) is a (k_1, \dots, k_μ) -tree of accepting transcripts for \mathbb{x} . Given \mathbb{x} and T as input, the special soundness extractor Ext outputs a valid witness for \mathbb{x} , contradicting the fact that \mathcal{D}_0 consists precisely of those (\mathbb{x}) such that $\mathbb{x} \notin L_R$. Finally, notice that by definition, for any $(\mathbb{x}, \tau) \in \text{PartTr}(\mu)$ with $(\mathbb{x}, \tau) \in \mathcal{D}$ and any $m \in \mathcal{M}_{\mu+1}$, we have that the verifier rejects (\mathbb{x}, τ, m) . This proves that Π is RBR sound with the claimed errors. \square

As a corollary of [Theorem 1.1](#) and [\[CMS19\]](#), we see that special sound IOPs can be compiled to secure non-interactive proofs in the quantum random oracle model via the BCS transformation [\[BCS16\]](#).

Corollary 1.6 (Special Soundness Implies FS Security in the QROM). *Let Π be a μ -round (k_1, \dots, k_μ) -special sound IOP. Let $\text{BCS}(\Pi)$ be the non-interactive proof obtained by applying the BCS transformation to Π in the random oracle model, and let $\varepsilon = \max_i \{(k_i - 1)/|C_i|\}$. Then $\text{BCS}(\Pi)$ has adaptive soundness error $O(t^2 \varepsilon + t^3/2^\lambda)$ against quantum attackers that make at most $t - O(q \log \ell)$ queries to the random oracle, where λ is the output length of the random oracle in bits, q is (an upper bound on) the total number of queries made by the verifier during any execution of Π , and ℓ is the total number of symbols sent by both the prover and verifier during any execution of Π .*

Proof. This is a direct consequence of [Theorem 1.1](#) and of [\[CMS19, Theorem 8.6\]](#). The former yields that any (k_1, \dots, k_μ) -special sound IOP has RBR soundness (in the usual, not generalized sense) at most

$\max_{i \in [\mu]} \{(k_i - 1)/|\mathcal{D}_i|\}$. The latter states that the BCS transformation of a RBR sound IOP is sound in the Quantum Random Oracle Model, with the parameters and soundness errors from the statement of the corollaries. \square

Next, we recall and prove [Theorem 1.2](#), which states that RBR knowledge implies special soundness.

Theorem 1.2 (RBR Knowledge Implies Special Soundness). *Let $\Pi = (P, V)$ be a μ -round IOP for a relation \mathbf{R} . Let C_i be the set of verifier challenges for round $i \in \{1, \dots, \mu\}$. Assume Π has round-by-round knowledge with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$, and let*

$$(k_1, \dots, k_\mu) = (\lceil |C_1| \varepsilon_1 \rceil + 1, \dots, \lceil |C_\mu| \varepsilon_\mu \rceil + 1).$$

Suppose $\sum_{i \in [\mu]} \prod_{j \in [i]} k_j$ is upper bounded by a polynomial (on the lengths of inputs). Then Π is (k_1, \dots, k_μ) -special sound.

Proof. Assume Π has RBR knowledge with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$. Let \mathcal{D} be a corresponding doomed set. Let $(k_1, \dots, k_\mu) = (\lceil |C_1| \varepsilon_1 \rceil + 1, \dots, \lceil |C_\mu| \varepsilon_\mu \rceil + 1)$, fix an input \mathbb{x} , and let T be a (k_1, \dots, k_μ) -tree of accepting transcripts for \mathbb{x} . For convenience we introduce the following notation. Let τ be a rooted path in T (i.e., a path starting at the root of T) with $(\mathbb{x}, \tau) \in \text{PartTr}(i)$ for some $i = 0, \dots, \mu$. Then by the definition of i -round partial transcripts, τ ends with a challenge c from C_i . As such, the path τ has a uniquely defined node after c in T . We denote it by $\text{EndNode}(\tau)$.

To complete the proof, the following claim is needed.

Claim 4.1. *There exists a rooted path τ in T with the following properties:*

- $(\mathbb{x}, \tau) \in \text{PartTr}(i - 1)$ and $(\mathbb{x}, \tau) \in \mathcal{D}$ for some $i \in [\mu]$; and
- for each of the k_i edges with challenges $(c_{i,j})_{j \in [k_i]}$ departing from node $\text{EndNode}(\tau)$, we have

$$(\mathbb{x}, \tau, \text{EndNode}(\tau), c_{i,j}) \notin \mathcal{D}.$$

Observe that for such a path τ satisfying [Claim 4.1](#), the transcript $(\mathbb{x}, \tau, \text{EndNode}(\tau))$ is RBR extractable since the fraction of challenges $c \in C_i$ such that $(\mathbb{x}, \tau, \text{EndNode}(\tau), c) \notin \mathcal{D}$ is at least $k_i/|C_i| > (k_i - 1)/|C_i| = \varepsilon_i$. Accordingly, we also call the path *RBR extractable*.

Assuming [Claim 4.1](#) is true, we can complete the proof of [Theorem 1.2](#) by defining a special soundness extractor Ext_{spec} for Π in the following way. Given an input \mathbb{x} and a (k_1, \dots, k_μ) -tree T' of accepting transcripts for \mathbb{x} , Ext_{spec} enumerates all paths in T' of the form (τ, m) where $(\mathbb{x}, \tau) \in \text{PartTr}(i)$ and $m = \text{EndNode}(\tau)$ for some $i = 0, \dots, \mu$. Then it runs the RBR knowledge extractor Ext_{rbr} on all the transcripts (\mathbb{x}, τ, m) . If Ext_{rbr} outputs a witness \mathbb{w} , then Ext_{spec} outputs the same \mathbb{w} , otherwise it outputs \perp (i.e., it aborts). By the above claim, we know that at least one of the transcripts (\mathbb{x}, τ, m) is RBR extractable so Ext_{spec} will output a valid witness and will not abort. Note that the total number of transcripts to inspect is exactly the number of vertices in T , which is exactly $1 + \sum_{i \in [\mu]} \prod_{j \in [i]} k_j$. By assumption, this is bounded by a polynomial in $|\mathbb{x}|$, and so Ext_{spec} runs in polynomial time. Thus the theorem follows, assuming our previous claim is true.

We now proceed with proving [Claim 4.1](#). For technical reasons, we actually prove the following more general claim.

Claim 4.2. *Suppose that τ is a rooted path in T and $(\mathbb{x}, \tau) \in \text{PartTr}(i - 1)$. Then there exists a RBR extractable rooted path in T of the form (τ, τ', m) , where $(\mathbb{x}, \tau, \tau') \in \text{PartTr}(j - 1)$ for some $j \geq i$ (if $j = i$, then $\tau' = \emptyset$), and $m = \text{EndNode}(\tau, \tau')$.*

Intuitively, [Claim 4.2](#) says that if we are given a rooted path τ in the tree (not necessarily a complete path), then the RBR extractable rooted path we are looking for can be found “by continuing τ ”. It is also easy to see that [Claim 4.2](#) implies [Claim 4.1](#).

We prove [Claim 4.2](#) by reverse induction on $i \in [\mu]$. Suppose $i = \mu$ and let (\mathbb{x}, τ) be as in the induction hypothesis. Let $c_{\mu,1}, \dots, c_{\mu,k_\mu}$ be all edges leaving from $m = \text{EndNode}(\tau)$. Notice that for all $t \in [k_\mu]$, the rooted path $(\tau, m, c_{\mu,t})$ is a μ -round partial transcript, and thus $(\mathbb{x}, \tau, m, c_{\mu,t}) \notin \mathcal{D}$. If this were not the case, the verifier would reject $(\mathbb{x}, \tau, m, c_{\mu,t}, m')$, where $m' = \text{EndNode}(\tau, m, c_{\mu,t}) \in \mathcal{M}_{\mu+1}$, contradicting the assumption that T is a tree of accepting transcripts. Hence $(\mathbb{x}, \tau, m, c_{\mu,t}) \notin \mathcal{D}$ for all $t \in [k_\mu]$.

Now fix $i \in [\mu - 1]$ and assume [Claim 4.2](#) holds for all j with $i < j \leq \mu$. Let $m = \text{EndNode}(\tau) \in \mathcal{M}_i$ and $c_{i,1}, \dots, c_{i,k_i}$'s be as before. If $(\mathbb{x}, \tau, m, c_{i,t}) \notin \mathcal{D}$ for all $t \in [k_i]$, then the claim is shown. Otherwise, there exists t_0 such that $(\mathbb{x}, \tau, m, c_{i,t_0}) \in \mathcal{D}$. Let $m' = \text{EndNode}(\mathbb{x}, \tau, m, c_{i,t_0}) \in \mathcal{M}_{i+1}$. Then $(\mathbb{x}, \tau, m, c_{i,t_0}, m')$ is of the form (\mathbb{x}, τ', m') , where $(\mathbb{x}, \tau') \in \text{PartTr}(i + 1)$ and $(\mathbb{x}, \tau') \in \mathcal{D}$. Applying our induction hypothesis on (\mathbb{x}, τ') , the claim then follows immediately. This completes the proof of [Claim 4.2](#) and thus completes the proof of the theorem. \square

Finally, we recall and prove [Theorem 1.3](#), which states that special soundness does not imply RBR knowledge.

Theorem 1.3 (Special Soundness does not Imply RBR Knowledge). *Assume $\text{NP} \neq \text{P}$. Then for any polynomial $\mu(|\mathbb{x}|)$, there exists a μ -round IOP Π with the following properties:*

- Π is $(1, \cdot^H, 1)$ -special sound.
- If Π is RBR knowledge sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$, then $\varepsilon_i(\ell) = 1$ for some input length ℓ and some $i \in \{1, \dots, \mu\}$.

Proof. Let \mathbf{R} be a relation such that $L_{\mathbf{R}}$ is a language in NP but not in P . Let $\mu = \mu(|\mathbb{x}|)$ be a polynomial. We define a μ -round IOP Π for \mathbf{R} as follows. Let $(\mathbb{x}, \mathbb{w}) \in \mathbf{R}$. For every round $i \in [\mu]$, the honest prover P sends a single bit 0 to the verifier V , and V responds with a uniformly sampled $c_i \stackrel{\$}{\leftarrow} C$. The last message sent by P is the witness \mathbb{w} . At the end of the protocol, V accepts only if $(\mathbb{x}, \mathbb{w}) \in \mathbf{R}$ and all messages sent by the prover are the single bit 0.

Clearly Π is sound with soundness error 0; moreover, Π is $(1, \cdot^H, 1)$ -special sound since any $(1, \cdot^H, 1)$ -tree of accepting transcripts for an input \mathbb{x} contains a valid witness \mathbb{w} for \mathbb{x} in its one leaf. Observe also that Π is perfectly complete, i.e., an honest prover convinces the verifier with probability 1.

Now assume Π has RBR knowledge with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$. Let \mathcal{D} be the corresponding doomed set of partial transcripts, and let Ext be the corresponding RBR knowledge extractor. We show the following claim.

Claim 4.3. *There exists $\ell \in \mathbb{N}$ such that for inputs \mathbb{x} of length ℓ and some $i \in [\mu]$, we have $\varepsilon_i(\ell) = 1$.*

Assume towards contradiction that this is not the case. Then there exists a deterministic polynomial time algorithm that given any $\mathbb{x} \in L_{\mathbf{R}}$ outputs a valid witness \mathbb{w} for \mathbb{x} . By our previous assumption, we have $\varepsilon_i(|\mathbb{x}|) < 1$ for all i . Let $i \in [\mu]$, let (\mathbb{x}, τ) be a $(i - 1)$ -round partial transcript with $(\mathbb{x}, \tau) \in \mathcal{D}$, and let $m \in \mathcal{M}_i$. Define

$$\rho(\mathbb{x}, \tau, m) := \Pr_{c \in C_i} [(\mathbb{x}, \tau, m, c) \notin \mathcal{D}].$$

Let (\mathbb{x}, τ') be a μ -round partial transcript generated from the interaction of the honest prover and the verifier. Then we have $(\mathbb{x}, \tau') = (\mathbb{x}, 0, c_1, 0, c_2, \dots, 0, c_\mu)$ for some challenges $c_i \stackrel{\$}{\leftarrow} C$. Observe that $(\mathbb{x}, \tau', \mathbb{w})$ is a complete transcript that is accepted by the verifier. Hence, by definition of RBR knowledge, $(\mathbb{x}, \tau') \notin \mathcal{D}$.

Let $\rho_{i-1} := \rho(\mathbb{x}, 0, c_1, \dots, 0, c_{i-1}, 0)$. We claim that $\rho_{i-1} = 1$ for some $i \in [\mu]$. If this was not the case, Π would not be perfectly complete since given (\mathbb{x}, \mathbb{w}) , the honest prover would only convince the verifier with probability at most $1 - \prod_{i \in [\mu]} (1 - \rho_{i-1})$ (i.e., the probability that, for some i , a doomed partial transcript $(\mathbb{x}, 0, c_1, \dots, 0, c_{i-1}, 0)$ escapes the doomed set after receiving a challenge c_i), but this latter probability is strictly less than 1 unless $\rho_{i-1} = 1$ for some $i \in [\mu]$. [Claim 4.3](#) follows.

Let $i \in [\mu]$ be such that $\rho_{i-1} = 1$. Then given $(\mathbb{x}, 0, c_1, \dots, 0, c_{i-1}, 0)$, the extractor Ext outputs a valid witness \mathbb{w}' for \mathbb{x} in polynomial time since $\rho_{i-1} > \varepsilon_i(|\mathbb{x}|)$. Notice that, conversely, if $\mathbb{x} \notin L_{\mathbf{R}}$, then given any partial transcript of the form $(\mathbb{x}, 0, c_1, \dots, c_{i-1}, 0)$, Ext outputs \perp in polynomial time.

This gives the following polynomial algorithm Ext' for the language $L_{\mathbf{R}}$. Given an arbitrary input \mathbb{x} , generate μ partial transcripts of the form $(\mathbb{x}, 0, c_1, \dots, 0, c_{i-1}, 0)$, for $i \in [\mu]$. Here the c_i are strings generated in an arbitrary deterministic manner. Give each of these as input to Ext . If for some of these Ext outputs a valid witness \mathbb{w}' for \mathbb{x} , then Ext' outputs \mathbb{w}' . Otherwise Ext' outputs \perp . The algorithm Ext' is deterministic and runs in polynomial time since (1) $\mu(|\mathbb{x}|)$ is polynomial; and (2) Ext runs in polynomial time. This contradicts the fact that $L_{\mathbf{R}}$ is in \mathbf{NP} but not in \mathbf{P} . Thus, our initial assumption that $\varepsilon_i(|\mathbb{x}|) < 1$ for all \mathbb{x} and all $i \in [\mu]$ cannot hold. \square

4.1 Special Unsoundness and Round-by-round Soundness are Dual

In this section we discuss relations between the soundness, the RBR soundness, and the special unsoundness of an IOP Π . We will see that RBR soundness and special unsoundness are “dual” concepts of each other, and that the former upper bounds the soundness of Π , while the latter lower bounds it, and moreover allows an attack to Π whose success probability is this lower bound. We will also see that when the RBR soundness and the special unsoundness errors are the same, then these errors are tight and their combination equals the soundness of Π .

We begin by formulating a slight variation of the definition of special unsoundness presented in [Definition 3.10](#). The main differences are that instead of having sets of “lucky challenges”, we have sets of “lucky partial transcripts”. The motivation behind this alternative formulation is that it highlights why special unsoundness is the dual notion of RBR soundness.

Definition 4.4 (Special Unsoundness [[AFK22](#)] – alternative formulation). *Let Π be a μ -round IOP for a relation \mathbf{R} . We say Π is special unsound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$ if there exist a set \mathcal{L} of “lucky” partial transcripts and an unbounded prover algorithm \mathbf{P}^* such that for all \mathbb{x} , the following hold.*

- For all $\mathbb{x} \notin L_{\mathbf{R}}$, we have that the 0-round partial transcript $(\mathbb{x}) \notin \mathcal{L}$.
- Let $(\mathbb{x}, \tau) \in \text{PartTr}(\mu)$ be a μ -round partial transcript. If $(\mathbb{x}, \tau) \in \mathcal{L}$, then $\forall (\mathbb{x}, \tau, m) = 1$ for any $m \in \mathcal{M}_{\mu+1}$. Moreover, for all $i \in [\mu - 1]$ and $(\mathbb{x}, \tau) \in \mathcal{L} \cap \text{PartTr}(i - 1)$, \mathbf{P}^* is able to compute $m \in \mathcal{M}_i$ such that $(\mathbb{x}, \tau, m, c) \in \mathcal{L}$ for all $c \in C_i$.
- Let $i \in [\mu]$, let $(\mathbb{x}, \tau) \in \text{PartTr}(i - 1)$ be a $(i - 1)$ -round partial transcript, and let $m \leftarrow \mathbf{P}^*(\mathbb{x}, \tau)$. Then

$$\Pr_{c \leftarrow C_i} [(\mathbb{x}, \tau, m, c) \in \mathcal{L}] \geq \varepsilon_i.$$

Remark 4.5 (Special Unsoundness as the Dual Notion of RBR Soundness). [Definition 4.4](#) can be understood as the dual of RBR soundness in the sense that the sets \mathcal{L}, \mathcal{D} of lucky and doomed transcripts from the respective definitions have opposite properties:

- An input $\mathbb{x} \notin \mathcal{L}$ does not belong to \mathcal{L} , while it belongs to \mathcal{D} .
- The verifier rejects any doomed complete transcript, while it accepts any complete lucky transcript. On the other hand, given a non-lucky partial transcript, if after some subsequent round the partial transcript is lucky, the verifier will eventually accept.
- For all round $i \in [\mu - 1]$, and for all prover, the probability that a $(i - 1)$ -round doomed partial transcripts stops being doomed at Round i is at most ε_i . On the other hand, there exists a prover that for all non-lucky $(i - 1)$ -round partial transcript, the probability that the transcript becomes lucky in the next round is at least ε'_i .

Remark 4.6 (Equivalence between [Definition 3.10](#) and [Definition 4.4](#)). [Definitions 3.10](#) and [4.4](#) are indeed equivalent: the set of “lucky” challenges from [Definition 3.10](#) depend on the partial transcript so far, hence one may as well define a set of “lucky” partial transcripts instead, as we did in [Definition 4.4](#). Moreover, the active and passive modes of the adversary P^* from [Definition 3.10](#) can be seen as analogues of P^* having or not produced a lucky partial transcript. If it has (passive mode), then P^* operates in a way that all subsequent transcripts are lucky, and so, following this analogy, P^* stays in passive mode until the end of the proof. As required, at the end of the proof, if the complete transcript is lucky (analogously, if P^* is in passive mode), the verifier accepts.

The next result relates the notions of soundness, RBR soundness, and special unsoundness of an IOP. The key observation is that the concepts of special unsoundness and round-by-round soundness are dual of each other. As a result, and intuitively speaking, we have that the special unsoundness error lower bounds the soundness of the protocol, while the round-by-round soundness error upper bounds it.

Theorem 1.4 (Relation between Soundness, Round-by-round Soundness, and Special Unsoundness). *Let Π be a μ -round IOP for a relation \mathbf{R} with soundness error ε . Then the following hold:*

1. **RBR soundness is an upper bound for soundness.** *If Π is round-by-round sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$, then*

$$\varepsilon \leq 1 - \prod_{i=1}^{\mu} (1 - \varepsilon_i)$$

for all $\mathbb{x} \notin L_{\mathbf{R}}$.

2. **Special unsoundness is a lower bound for soundness.** *If Π is special unsound with errors $(\varepsilon'_1, \dots, \varepsilon'_\mu)$, then*

$$1 - \prod_{i=1}^{\mu} (1 - \varepsilon'_i) \leq \varepsilon$$

for all $\mathbb{x} \notin L_{\mathbf{R}}$. Moreover, there exists a dishonest unbounded prover P^ that, given any input \mathbb{x} , manages to make the verifier accept with probability at least $1 - \prod_i (1 - \varepsilon'_i)$.*

3. **Tightness of RBR soundness, soundness, and special unsoundness.** *Suppose Π is round-by-round sound with errors $(\varepsilon_1, \dots, \varepsilon_\mu)$ and that Π is special unsound with the same errors $(\varepsilon_1, \dots, \varepsilon_\mu)$. Then*

$$\varepsilon = 1 - \prod_{i=1}^{\mu} (1 - \varepsilon_i).$$

Moreover, the error is tight in the sense that there exists a dishonest prover P^ that, given any input \mathbb{x} , manages to have the verifier accept with probability at least ε .*

Remark 4.7. The quantity $\rho := 1 - \prod_{i \in [\mu]} (1 - x_i)$ is, for small x_i 's, approximately $\sum_{i \in [\mu]} x_i$, as this is the first-order term in the Taylor approximation of ρ around the point $(0, \dots, 0)$.

Proof. We begin by proving Item 1. Assume $\mathbb{x} \notin L_{\mathbf{R}}$ and let P^* be any unbounded dishonest prover. Let $(\mathbb{x}, \tau) \in \text{PartTr}(\mu)$ be a μ -round partial transcript produced during the interaction of P^* and V , and let $m \in \mathcal{M}_\mu$ so that (\mathbb{x}, τ, m) is a complete transcript. Let $\rho := 1 - \prod_{i \in [\mu]} (1 - \varepsilon_i)$.

Define $p := \Pr_{(\tau, m) \leftarrow \langle P^*, V \rangle(\mathbb{x})} [V(\mathbb{x}, \tau) = 1]$. Then $p \leq \rho$ for all input length ℓ and all (\mathbb{x}, τ) , with $|\mathbb{x}| = \ell$. To ensure $V(\mathbb{x}, \tau, m) = 1$, it is necessary (but not sufficient) that $(\mathbb{x}, \tau) \notin \mathcal{D}_\mu$. Hence

$$\Pr_{(\tau, m) \leftarrow \langle P^*, V \rangle(\mathbb{x})} [V(\mathbb{x}, \tau) = 1] \leq \Pr_{(\tau, m) \leftarrow \langle P^*, V \rangle(\mathbb{x})} [(\mathbb{x}, \tau) \notin \mathcal{D}_\mu].$$

Let E_i be the event that the i -round partial transcript (\mathbb{x}, τ) is not in the doomed set \mathcal{D} . Assume that if event E_i occurs, then P^* can send messages m_j for any $i + 1 \leq j \leq \mu$ to ensure that events E_{i+1}, \dots, E_μ each occur with probability 1. Then $\Pr_{(\tau, m) \leftarrow \langle P^*, V \rangle(\mathbb{x})} [(\mathbb{x}, \tau) \notin \mathcal{D}_\mu]$ is exactly the probability that there exists a round $i \in [\mu]$ such that event E_i occurs (i.e., P^* can escape the doomed set in round i). This follows by our assumption that if E_i occurs then P^* can send messages such that E_{i+1}, \dots, E_μ each happen with probability 1. By the assumption of RBR soundness, E_i occurs with probability at most ε_i , and hence we have

$$\Pr_{(\tau, m) \leftarrow \langle P^*, V \rangle(\mathbb{x})} [(\mathbb{x}, \tau) \notin \mathcal{D}_\mu] \leq \rho = 1 - \prod_{i \in [\mu]} (1 - \varepsilon_i). \quad (3)$$

Now consider any other attacker $P^{*'} such that once it has obtained some a partial transcript $(\mathbb{x}, \tau) \notin \mathcal{D}$, it produces a transcript extension τ' such that $\Pr_c [(\mathbb{x}, \tau, \tau', c) \notin \mathcal{D}] > 0$. We claim this attacker always produces a non-doomed complete transcript $(\mathbb{x}, \tau) \notin \mathcal{D}$ with at most the probability that the previous attacker P^* does. This is because once a previously doomed partial transcript has been extended so that it is not in a doomed set, P^* ends with a non-doomed complete transcript with probability 1, while the probability of $P^{*'}$ may be smaller. This completes the proof of Item 1.$

Next we prove Item 2. Fix an input $\mathbb{x} \notin L_{\mathbf{R}_\mathbb{x}}$. Let P^* be the dishonest prover from [Definition 4.4](#). As before, let $(\mathbb{x}, \tau) \in \text{PartTr}(\mu)$ be a μ -round transcript produced during the interaction of the prover and the verifier, and let $m \in \mathcal{M}_\mu$. In order to have $V(\mathbb{x}, \tau, m) = 1$, it is sufficient (but not necessary) that there exists a prefix τ_i of τ such that (\mathbb{x}, τ_i) is an i -round partial transcript and $(\mathbb{x}, \tau_i) \in \mathcal{L}$. Let E be the event that such a prefix exists. Similar to Item 1, the probability of E occurring is at least $\rho' := 1 - \prod_{i \in [\mu]} (1 - \varepsilon'_i)$; i.e., the probability that after μ trials, each with success probability ε'_i , at least one trial was successful. Hence

$$\Pr_{(\tau, m) \leftarrow \langle P^*, V \rangle(\mathbb{x})} [V(\mathbb{x}, \tau, m) = \text{accept}] \geq \rho', \quad (4)$$

proving Item 2.

Item 3 follows directly from Items 1 and 2, since we have that any malicious prover convinces the verifier with probability at most $1 - \prod_{i \in [\mu]} (1 - \varepsilon_i)$, and that there exists a prover that convinces the verifier with at least this probability. \square

Remark 4.8. In [\[CCH⁺18\]](#), the authors prove that, given a μ -round interactive proof/argument Π with soundness $\varepsilon_{\text{sound}}$ and round-by-round soundness ε_{rbr} (in the non-generalized sense of RBR soundness), one has $\varepsilon_{\text{sound}} \leq \mu \varepsilon_{\text{rbr}}$. The previous [Theorem 1.4](#) yields a slight improvement in this formula, giving $\varepsilon_{\text{sound}} \leq 1 - (1 - \varepsilon_{\text{rbr}})^\mu$. We remark however that we expect this improvement to be known already (at least in folklore), and that when the errors are small, the improvement is negligible (see [Remark 4.7](#)).

Acknowledgements

Alexander R. Block was supported by DARPA under Contract No. HR00112020022 and No. HR00112020025. Albert Garreta was supported by the Ethereum Foundation’s grant FY23-0885. Pratyush Ranjan Tiwari was partly supported by the Ethereum Foundation’s grant FY23-1087, a Security & Privacy research gift from Google, and a research gift from Cisco. Michał Zając was supported by the Ethereum Foundation’s grant FY23-0885. The views, opinions, findings, conclusions and/or recommendations expressed in this material are those of the authors and should not be interpreted as reflecting the position or policy of DARPA or the United States Government, and no official endorsement should be inferred.

References

- [ACF21] Thomas Attema, Ronald Cramer, and Serge Fehr. Compressing proofs of k-out-of-n partial knowledge. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part IV*, volume 12828 of *LNCS*, pages 65–91, Virtual Event, August 2021. Springer, Heidelberg. [doi:10.1007/978-3-030-84259-8_3](https://doi.org/10.1007/978-3-030-84259-8_3). 3
- [ACK21] Thomas Attema, Ronald Cramer, and Lisa Kohl. A compressed Σ -protocol theory for lattices. In Tal Malkin and Chris Peikert, editors, *CRYPTO 2021, Part II*, volume 12826 of *LNCS*, pages 549–579, Virtual Event, August 2021. Springer, Heidelberg. [doi:10.1007/978-3-030-84245-1_19](https://doi.org/10.1007/978-3-030-84245-1_19). 7
- [AFK22] Thomas Attema, Serge Fehr, and Michael Klooß. Fiat-shamir transformation of multi-round interactive proofs. In Eike Kiltz and Vinod Vaikuntanathan, editors, *Theory of Cryptography - 20th International Conference, TCC 2022, Chicago, IL, USA, November 7-10, 2022, Proceedings, Part I*, 2022. 2, 3, 6, 7, 12, 16, 21
- [Bab85] László Babai. Trading group theory for randomness. In Robert Sedgewick, editor, *Proceedings of the 17th Annual ACM Symposium on Theory of Computing, May 6-8, 1985, Providence, Rhode Island, USA*, pages 421–429. ACM, 1985. 2
- [BBHR18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast reed-solomon interactive oracle proofs of proximity. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, 2018. 1
- [BCGT13] Eli Ben-Sasson, Alessandro Chiesa, Daniel Genkin, and Eran Tromer. Fast reductions from RAMs to delegatable succinct constraint satisfaction problems: extended abstract. In Robert D. Kleinberg, editor, *ITCS 2013*, pages 401–414. ACM, January 2013. [doi:10.1145/2422436.2422481](https://doi.org/10.1145/2422436.2422481). 7
- [BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In Martin Hirt and Adam D. Smith, editors, *TCC 2016-B, Part II*, volume 9986 of *LNCS*, pages 31–60. Springer, Heidelberg, October / November 2016. [doi:10.1007/978-3-662-53644-5_2](https://doi.org/10.1007/978-3-662-53644-5_2). 1, 2, 4, 5, 6, 7, 10, 18
- [BDG⁺13] Nir Bitansky, Dana Dachman-Soled, Sanjam Garg, Abhishek Jain, Yael Tauman Kalai, Adriana López-Alt, and Daniel Wichs. Why “Fiat-Shamir for proofs” lacks a proof. In Amit Sahai, editor, *TCC 2013*, volume 7785 of *LNCS*, pages 182–201. Springer, Heidelberg, March 2013. [doi:10.1007/978-3-642-36594-2_11](https://doi.org/10.1007/978-3-642-36594-2_11). 7

- [BGKS20] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, and Shubhangi Saraf. DEEP-FRI: sampling outside the box improves soundness. In *Innovations in Theoretical Computer Science Conference, ITCS*, 2020. 1
- [BTVW14] Andrew J. Blumberg, Justin Thaler, Victor Vu, and Michael Walfish. Verifiable computation using multiple provers. Cryptology ePrint Archive, Report 2014/846, 2014. <https://eprint.iacr.org/2014/846>. 7
- [CBBZ23] Binyi Chen, Benedikt Bünz, Dan Boneh, and Zhenfei Zhang. Hyperplonk: Plonk with linear-time prover and high-degree custom gates. In *Advances in Cryptology - EUROCRYPT*, 2023. 1
- [CCH⁺18] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, and Ron D. Rothblum. Fiat-Shamir from simpler assumptions. Cryptology ePrint Archive, Report 2018/1004, 2018. <https://eprint.iacr.org/2018/1004>. 23
- [CCH⁺19] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs. Fiat-Shamir: from practice to theory. In Moses Charikar and Edith Cohen, editors, *51st ACM STOC*, pages 1082–1090. ACM Press, June 2019. doi:10.1145/3313276.3316380. 2, 6, 7, 16, 17
- [CCRR18] Ran Canetti, Yilei Chen, Leonid Reyzin, and Ron D. Rothblum. Fiat-Shamir and correlation intractability from strong KDM-secure encryption. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 91–122. Springer, Heidelberg, April / May 2018. doi:10.1007/978-3-319-78381-9_4. 7
- [CDS94] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Yvo Desmedt, editor, *Advances in Cryptology - CRYPTO '94, 14th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1994, Proceedings*, volume 839 of *Lecture Notes in Computer Science*, pages 174–187. Springer, 1994. 2
- [CMS19] Alessandro Chiesa, Peter Manohar, and Nicholas Spooner. Succinct arguments in the quantum random oracle model. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 1–29. Springer, Heidelberg, December 2019. doi:10.1007/978-3-030-36033-7_1. 1, 2, 4, 5, 6, 10, 16, 17, 18
- [CMSZ21] Alessandro Chiesa, Fermi Ma, Nicholas Spooner, and Mark Zhandry. Post-quantum succinct arguments: Breaking the quantum rewinding barrier. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 49–58. IEEE, 2021. doi:10.1109/FOCS52979.2021.00014. 5
- [CMT12] Graham Cormode, Michael Mitzenmacher, and Justin Thaler. Practical verified computation with streaming interactive proofs. In Shafi Goldwasser, editor, *ITCS 2012*, pages 90–112. ACM, January 2012. doi:10.1145/2090236.2090245. 7
- [COS20] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. Fractal: Post-quantum and transparent recursive proofs from holography. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part I*, volume 12105 of *LNCS*, pages 769–793. Springer, Heidelberg, May 2020. doi:10.1007/978-3-030-45721-1_27. 1, 2, 4, 6, 7

- [DFMS19] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 356–383. Springer, Heidelberg, August 2019. doi:10.1007/978-3-030-26951-7_13. 5
- [FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO’86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987. doi:10.1007/3-540-47721-7_12. 2, 7
- [GKR08] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 113–122. ACM Press, May 2008. doi:10.1145/1374376.1374396. 7
- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 1989. 2
- [GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. Cryptology ePrint Archive, Report 2019/953, 2019. <https://eprint.iacr.org/2019/953>. 1
- [HL18] Justin Holmgren and Alex Lombardi. Cryptographic hashing from strong one-way functions (or: One-way product functions and their applications). In Mikkel Thorup, editor, *59th FOCS*, pages 850–858. IEEE Computer Society Press, October 2018. doi:10.1109/FOCS.2018.00085. 7
- [Hol19] Justin Holmgren. On round-by-round soundness and state restoration attacks. Cryptology ePrint Archive, Report 2019/1261, 2019. <https://eprint.iacr.org/2019/1261>. 2, 4, 6, 7, 16
- [Kil92] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *24th ACM STOC*, pages 723–732. ACM Press, May 1992. doi:10.1145/129712.129782. 7
- [KPV19] Assimakis Kattis, Konstantin Panarin, and Alexander Vlasov. Redshift: Transparent snarks from list polynomial commitments. Cryptology ePrint Archive, Paper 2019/1400, 2019. <https://eprint.iacr.org/2019/1400>. URL: <https://eprint.iacr.org/2019/1400>, doi:10.1145/548606.3560657. 1, 2, 3, 4
- [KRR17] Yael Tauman Kalai, Guy N. Rothblum, and Ron D. Rothblum. From obfuscation to the security of Fiat-Shamir for proofs. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 224–251. Springer, Heidelberg, August 2017. doi:10.1007/978-3-319-63715-0_8. 7
- [LZ19] Qipeng Liu and Mark Zhandry. Revisiting post-quantum Fiat-Shamir. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 326–355. Springer, Heidelberg, August 2019. doi:10.1007/978-3-030-26951-7_12. 5
- [Mic94] Silvio Micali. CS proofs (extended abstracts). In *35th FOCS*, pages 436–453. IEEE Computer Society Press, November 1994. doi:10.1109/SFCS.1994.365746. 7
- [Pol22] Polygon Zero Team. Plonky2: Fast recursive arguments with plonk and fri, 2022. <https://github.com/mir-protocol/plonky2/tree/main/plonky2>. 1

- [RR20] Noga Ron-Zewi and Ron D. Rothblum. Local proofs approaching the witness length [extended abstract]. In *61st FOCS*, pages 846–857. IEEE Computer Society Press, November 2020. doi:10.1109/FOCS46700.2020.00083. 7
- [RRR21] Omer Reingold, Guy N. Rothblum, and Ron D. Rothblum. Constant-round interactive proofs for delegating computation. *SIAM J. Comput.*, 2021. 1, 7
- [Set20] Srinath Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 704–737. Springer, Heidelberg, August 2020. doi:10.1007/978-3-030-56877-1_25. 7
- [Tha13] Justin Thaler. Time-optimal interactive proofs for circuit evaluation. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part II*, volume 8043 of *LNCS*, pages 71–89. Springer, Heidelberg, August 2013. doi:10.1007/978-3-642-40084-1_5. 7
- [Wik18] Douglas Wikström. Special soundness revisited. Cryptology ePrint Archive, Report 2018/1157, 2018. <https://eprint.iacr.org/2018/1157>. 2
- [Wik21] Douglas Wikström. Special soundness in the random oracle model. Cryptology ePrint Archive, Report 2021/1265, 2021. <https://eprint.iacr.org/2021/1265>. 2
- [WTs⁺18] Riad S. Wahby, Ioanna Tzialla, abhi shelat, Justin Thaler, and Michael Walfish. Doubly-efficient zkSNARKs without trusted setup. In *2018 IEEE Symposium on Security and Privacy*, pages 926–943. IEEE Computer Society Press, May 2018. doi:10.1109/SP.2018.00060. 7