



Coin Flipping of *Any* Constant Bias Implies One-Way Functions

Itay Berman^{*§}Iftach Haitner^{†§}Aris Tentes^{‡§}

June 21, 2015

Abstract

We show that the existence of a coin-flipping protocol safe against *any* non-trivial constant bias (e.g., .499) implies the existence of one-way functions. This improves upon a recent result of [Haitner and Omri](#) [FOCS '11], who proved this implication for protocols with bias $\frac{\sqrt{2}-1}{2} - o(1) \approx .207$. Unlike the result of [Haitner and Omri](#), our result also holds for *weak* coin-flipping protocols.

Keywords: coin-flipping protocols; one-way functions; minimal hardness assumptions

*MIT Computer Science and Artificial Intelligence Laboratory. E-mail: itayberm@mit.edu. Most of this work was done while the author was in the School of Computer Science, Tel Aviv University.

†School of Computer Science, Tel Aviv University. E-mail: iftachh@cs.tau.ac.il.

‡E-mail: aristent@gmail.com. Most of this work was done while the author was in the Department of Computer Science, New York University.

§Research supported by ISF grant 1076/11, the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11), US-Israel BSF grant 2010196 and Check Point Institute for Information Security.

Contents

1	Introduction	1
1.1	Our Result	1
1.2	Related Results	2
1.3	Our Techniques	2
1.4	Open Questions	9
2	Preliminaries	9
2.1	Notations	9
2.2	Two-Party Protocols	10
2.3	Coin-Flipping Protocols	12
2.4	One-Way Functions and Distributional One-Way Functions	13
2.5	Two Inequalities	15
3	The Biased-Continuation Attack	15
3.1	Basic Observations About $A^{(i)}$	17
3.2	Optimal Valid Attacks	19
3.3	Dominated Measures	20
3.4	Warmup — Proof Attempt Using a (Single) Dominated Measure	24
3.5	Back to the Proof — Sequence of Alternating Dominated Measures	27
3.6	Improved Analysis Using Alternating Dominated Measures	33
3.7	Proving Lemma 3.25	35
3.8	Proving Lemma 3.26	46
4	Efficiently Biasing Coin-Flipping Protocols	48
4.1	The Approximated Biased-Continuation Attacker	49
4.2	Attacking Pruned Protocols	58
4.3	The Pruning-in-the-Head Attacker	74
4.4	Main Theorem — Constructing an Efficient Attacker	82
A	Missing Proofs	88
A.1	Proving Lemma 2.17	88
A.2	Proving Lemma 2.18	89

1 Introduction

A central focus of modern cryptography has been to investigate the weakest possible assumptions under which various cryptographic primitives exist. This direction of research has been quite fruitful, and minimal assumptions are known for a wide variety of primitives. In particular, it has been shown that one-way functions (i.e., easy to compute but hard to invert) imply pseudorandom generators, pseudorandom functions, symmetric-key encryption/message authentication, commitment schemes, and digital signatures [9, 10, 13, 12, 20, 21, 8, 23], where one-way functions were also shown to be implied by each of these primitives [15].

An important exception to the above successful characterization is that of coin-flipping (-tossing) protocols. A coin-flipping protocol [3] allows the honest parties to jointly flip an unbiased coin, where even a cheating (efficient) party cannot bias the outcome of the protocol by very much. Specifically, a coin-flipping protocol is δ -bias if no efficient cheating party can make the common output to be 1, or to be 0, with probability greater than $\frac{1}{2} + \delta$. While one-way functions are known to imply negligible-bias coin-flipping protocols [3, 20, 13], the other direction is less clear. Impagliazzo and Luby [15] showed that $\Theta(1/\sqrt{m})$ -bias coin-flipping protocols imply one-way functions, where m is the number of rounds in the protocol.¹ Recently, Maji, Prabhakaran, and Sahai [17] extended the above for $(\frac{1}{2} - 1/\text{poly}(n))$ -bias *constant-round* protocols, where n is the security parameter. More recently, Haitner and Omri [11] showed that the above implication holds for $(\frac{\sqrt{2}-1}{2} - o(1) \approx 0.207)$ -bias coin-flipping protocols (of arbitrary round complexity). No such implications were known for any other choice of parameters, and in particular for protocols with bias greater than $\frac{\sqrt{2}-1}{2}$ with super-constant round complexity.

1.1 Our Result

In this work, we make progress towards answering the question of whether coin-flipping protocols also imply one-way functions. We show that (even weak) coin-flipping protocols, safe against any non-trivial bias (e.g., 0.4999), do in fact imply such functions. We note that unlike [11], but like [15, 17], our result also applies to the so-called *weak coin-flipping protocols* (see Section 2.3 for the formal definition of strong and weak coin-flipping protocols). Specifically, we prove the following theorem.

Theorem 1.1 (informal). *For any $c > 0$, the existence of a $(\frac{1}{2} - c)$ -bias coin-flipping protocol (of any round complexity) implies the existence of one-way functions.*

Note that $\frac{1}{2}$ -bias coin-flipping protocol requires no assumption (i.e., one party flips a coin and announces the result to the other party). So our result is tight as long as constant biases (i.e., independent of the security parameter) are involved.

To prove Theorem 1.1, we observe a connection between the success probability of the best (valid) attacks in a two-party game (i.e., chess) and the success of the biased-continuation attack of [11] in winning this game (see more in Section 1.3). The implications of this interesting connection seem to extend beyond the question at the focus of this paper.

¹In [15], only $\text{neg}(m)$ -bias was stated. Proving the same implication for $\Theta(1/\sqrt{m})$ -bias follows from the proof outlined in [15] and the result by Cleve and Impagliazzo [6].

1.2 Related Results

As mentioned above, Impagliazzo and Luby [15] showed that negligible-bias coin-flipping protocols imply one-way functions. Maji et al. [17] proved the same for $(\frac{1}{2} - o(1))$ -bias yet constant-round protocols. Finally, Haitner and Omri [11] showed that the above implication holds for $\frac{\sqrt{2}-1}{2} - o(1) \approx 0.207$ -bias (strong) coin-flipping protocols (of arbitrary round complexity). Results of weaker complexity implications are also known.

Zachos [24] has shown that non-trivial (i.e., $(\frac{1}{2} - o(1))$ -bias), constant-round coin-flipping protocols imply that $\text{NP} \not\subseteq \text{BPP}$, where Maji et al. [17] proved the same implication for $(\frac{1}{4} - o(1))$ -bias coin-flipping protocols of arbitrary round complexity. Finally, it is well known that the existence of non-trivial coin-flipping protocols implies that $\text{PSPACE} \not\subseteq \text{BPP}$. Apart from [11], all the above results extend to weak coin-flipping protocols. See Table 1 for a summary.

<i>Implication</i>	<i>Protocol type</i>	<i>Paper</i>
Existence of OWFs	$(\frac{1}{2} - c)$ -bias, for some $c > 0$	This work
Existence of OWFs	$(\frac{\sqrt{2}-1}{2} - o(1))$ -bias	Haitner and Omri [11] ²
Existence of OWFs	$(\frac{1}{2} - o(1))$ -bias, <i>constant round</i>	Maji et al. [17]
Existence of OWFs	Negligible bias	Impagliazzo and Luby [15]
$\text{NP} \not\subseteq \text{BPP}$	$(\frac{1}{4} - o(1))$ -bias	Maji et al. [17]
$\text{NP} \not\subseteq \text{BPP}$	$(\frac{1}{2} - o(1))$ -bias, <i>constant round</i>	Zachos [24]
$\text{PSPACE} \not\subseteq \text{BPP}$	Non-trivial	Common knowledge

Table 1: Results summary.

Information theoretic coin-flipping protocols (i.e., whose security holds against all-powerful attackers) were shown to exist in the quantum world; Mochon [18] presented an ε -bias quantum weak coin-flipping protocol for any $\varepsilon > 0$. Chailloux and Kerenidis [4] presented a $(\frac{\sqrt{2}-1}{2} - \varepsilon)$ -bias quantum strong coin-flipping protocol for any $\varepsilon > 0$ (this bias was shown in [16] to be tight). A key step in [4] is a reduction from strong to weak coin-flipping protocols, which holds also in the classical world.

A related line of work considers *fair* coin-flipping protocols. In this setting the honest party is required to always output a bit, whatever the other party does. In particular, a cheating party might bias the output coin just by aborting. We know that one-way functions imply fair $(1/\sqrt{m})$ -bias coin-flipping protocols [1, 5], where m is the round complexity of the protocol, and this quantity is known to be tight for $O(m/\log m)$ -round protocols with fully black-box reductions [7]. Oblivious transfer, on the other hand, implies fair $1/m$ -bias protocols [19, 2] (this bias was shown in [5] to be tight).

1.3 Our Techniques

The following is a rather elaborate, high-level description of the ideas underlying our proof.

That the existence of a given (cryptographic) primitive implies the existence of one-way functions is typically proven by looking at the *primitive core function* — an efficiently computable

¹Only holds for *strong* coin-flipping protocols.

function (not necessarily unique) whose inversion on uniformly chosen outputs implies breaking the security of the primitive.³ For private-key encryption, for instance, a possible core function is the mapping from the inputs of the encryption algorithm (i.e., message, secret key, and randomness) into the ciphertexts. Assuming that one has defined such a core function for a given primitive, then, by definition, this function should be one-way. So it all boils down to finding, or proving the existence of, such a core function for the primitive under consideration. For a *non-interactive* primitive, finding such a core function is typically easy. In contrast, for an *interactive* primitive, finding such a core function, or functions is, at least in many settings, a much more involved task. The reason is that in order to break an interactive primitive, the attacker typically has to invert a given function on many different outputs, where these outputs are chosen *adaptively* by the attacker, after seeing the answers to the previous queries. As a result, it is very challenging to find a single function, or even finitely many functions, whose output distribution (on uniformly chosen input) matches the distribution of the attacker’s queries.⁴

The only plausible candidate to serve as the core function of a coin-flipping protocol would seem to be its *transcript function*: the function that maps the parties’ randomness into the resulting protocol transcript (i.e., the transcript produced by executing the protocol with this randomness). In order to bias the output of an m -round coin-flipping protocol by more than $O(\frac{1}{\sqrt{m}})$, a super-constant number of adaptive inversions of the transcript function seems necessary. Yet we managed to prove that the transcript function is the core function of any (constant-bias) coin-flipping protocol. This is done by designing an adaptive attacker for any such protocol whose query distribution is “not too far” from the output distribution of the transcript function (when invoked on uniform inputs). Since our attacker, described below, is not only adaptive, but also defined in a recursive manner, proving that it possesses the aforementioned property was one of the major challenges we faced.

In what follows, we give a high-level overview of our attacker that ignores computational issues (i.e., assumes it has a perfect inverter for any function). We then explain how to adjust this attacker to work with the inverter of the protocol’s transcript function.

1.3.1 Optimal Valid Attacks and The Biased-Continuation Attack

The crux of our approach lies in an interesting connection between the optimal attack on a coin-flipping protocol and the more feasible, *recursive biased-continuation* attack. The latter attack recursively applies the biased-continuation attack used by Haitner and Omri [11] to achieve their constant-bias attack (called there, the *random-continuation* attack) and is the basis of our efficient attack (assuming one-way functions do not exist) on coin-flipping protocols. The results outlining the aforementioned connection, informally stated in this section and formally stated and proven in Section 3, hold for any two-player full information game with binary common outcome.

Let $\Pi = (\mathbf{A}, \mathbf{B})$ be a coin-flipping protocol (i.e., the common output of the honest parties is a uniformly chosen bit). In this discussion we restrict ourselves to analyzing attacks that, when carried out by the left-hand party, i.e., \mathbf{A} , are used to bias the outcome towards one, and when

³For the sake of this informal discussion, inverting a function on a given value means returning a *uniformly* chosen preimage of this value.

⁴If the attacker makes a *constant* number of queries, one can overcome the above difficulty by defining a set of core functions f_1, \dots, f_k , where f_1 is the function defined by the primitive, f_2 is the function defined by the attacker after making the first inversion call, and so on. Since the evaluation time of f_{i+1} is polynomial in the evaluation time of f_i (since evaluating f_{i+1} requires a call to an inverter of f_i), this approach fails miserably for attackers of super-constant query complexity.

carried out by the right-hand party, i.e., \mathcal{B} , are used to bias the outcome towards zero. Analogous statements hold for opposite attacks (i.e., attacks carried out by \mathcal{A} and used to bias towards zero, and attacks carried out by \mathcal{B} and used to bias towards one). The optimal valid attacker \mathcal{A} carries out the *best* attack \mathcal{A} can employ (using unbounded power) to bias the protocol towards *one*, while sending *valid* messages — ones that could have been sent by the honest party. The optimal valid attacker \mathcal{B} , carrying out the best attack \mathcal{B} can employ to bias the protocol towards *zero*, is analogously defined. Since, without loss of generality, the optimal valid attackers are deterministic, the expected outcome of $(\mathcal{A}, \mathcal{B})$ is either zero or one. As a first step, we give a lower bound on the success probability of the recursive biased-continuation attack carried out by the party winning the aforementioned game. As this lower bound might not be sufficient for our goal (it might be less than constant) — and this is a crucial point in the description below — our analysis takes additional steps to give an arbitrarily-close-to-one lower bound on the success probability of the recursive biased-continuation attack carried out by *some* party, which may or may not be the same party winning the aforementioned game.⁵

Assume that \mathcal{A} is the winning party when playing against \mathcal{B} . Since \mathcal{A} sends only valid messages, it follows that the expected outcome of $(\mathcal{A}, \mathcal{B})$, i.e., honest \mathcal{A} against the optimal attacker for \mathcal{B} , is larger than zero (since \mathcal{A} might send the optimal messages “by mistake”). Let $\text{OPT}_{\mathcal{A}}(\Pi)$ be the expected outcome of the protocol $(\mathcal{A}, \mathcal{B})$ and let $\text{OPT}_{\mathcal{B}}(\Pi)$ be 1 minus the expected outcome of the protocol $(\mathcal{A}, \mathcal{B})$. The above observation yields that $\text{OPT}_{\mathcal{A}}(\Pi) = 1$, while $\text{OPT}_{\mathcal{B}}(\Pi) = 1 - \alpha < 1$. This gives rise to the following question: *what gives \mathcal{A} an advantage over \mathcal{B} ?*

We show that if $\text{OPT}_{\mathcal{B}}(\Pi) = 1 - \alpha$, then there exists an α -dense set $\mathcal{S}^{\mathcal{A}}$ of 1-transcripts, full transcripts in which the parties’ common output is 1,⁶ that are “dominated by \mathcal{A} ”. The \mathcal{A} -dominated set has an important property — its density is “immune” to any action \mathcal{B} might take, even if \mathcal{B} is employing its optimal attack; specifically, the following holds:

$$\Pr_{\langle \mathcal{A}, \mathcal{B} \rangle} [\mathcal{S}^{\mathcal{A}}] = \Pr_{\langle \mathcal{A}, \mathcal{B} \rangle} [\mathcal{S}^{\mathcal{A}}] = \alpha, \quad (1)$$

where $\langle \Pi' \rangle$ samples a random full transcript of protocol Π' . It is easy to see that the above holds if \mathcal{A} controls the root of the tree and has a 1-transcript as a direct descendant; see Figure 1 for a concrete example. The proof of the general case can be found in Section 3. Since the \mathcal{A} -dominated set is \mathcal{B} -immune, a possible attack for \mathcal{A} is to go towards this set. Hence, what seems like a feasible adversarial attack for \mathcal{A} is to mimic \mathcal{A} ’s attack by hitting the \mathcal{A} -dominated set with high probability. It turns out that the biased-continuation attack of [11] does exactly that.

The biased-continuation attacker $\mathcal{A}^{(1)}$, taking the role of \mathcal{A} in Π and trying to bias the output of Π towards one, is defined as follows: given that the partial transcript is trans , algorithm $\mathcal{A}^{(1)}$ samples a pair of random coins $(r_{\mathcal{A}}, r_{\mathcal{B}})$ that is consistent with trans and leads to a 1-transcript, and then acts as the honest \mathcal{A} on the random coins $r_{\mathcal{A}}$, given the transcript trans . In other words, $\mathcal{A}^{(1)}$ takes the first step of a random continuation of $(\mathcal{A}, \mathcal{B})$ leading to a 1-transcript. (The attacker $\mathcal{B}^{(1)}$, taking the role of \mathcal{B} and trying to bias the outcome towards zero, is analogously defined.) Haitner and Omri [11] showed that for any coin-flipping protocol, if either \mathcal{A} or \mathcal{B} carries out the biased-continuation attack towards one, the outcome of the protocol will be biased towards one by $\frac{\sqrt{2}-1}{2}$

⁵That the identity of the winner in $(\mathcal{A}, \mathcal{B})$ cannot be determined by the recursive biased-continuation attack is crucial. Since we show that the latter attack can be efficiently approximated assuming one-way functions do not exist, the consequences of giving up this information would be profound. It would mean that we can estimate the optimal attack (which is implemented in PSPACE) using only the assumption that one-way functions do not exist.

⁶Throughout, we assume without loss of generality that the protocol’s transcript determines the common output of the parties.

(when interacting with the honest party).⁷ Our basic attack employs the above biased-continuation attack recursively. Specifically, for $i > 1$ we consider the attacker $A^{(i)}$ that takes the first step of a random continuation of $(A^{(i-1)}, B)$ leading to a 1-transcript, letting $A^{(0)} \equiv A$. The attacker $B^{(i)}$ is analogously defined. Our analysis takes a different route from that of [11], whose approach is only applicable for handling bias up to $\frac{\sqrt{2}-1}{2}$ and cannot be applied to weak coin-flipping protocols.⁸ Instead, we analyze the probability of the biased-continuation attacker to hit the dominated set we introduced above.

Let trans be a 1-transcript of Π in which all messages are sent by A . Since $A^{(1)}$ picks a random 1-transcript, and B cannot force $A^{(1)}$ to diverge from this transcript, the probability to produce trans under an execution of $(A^{(1)}, B)$ is *doubled* with respect to this probability under an execution of (A, B) (assuming the expected outcome of (A, B) is $1/2$). The above property, that B cannot force $A^{(1)}$ to diverge from a transcript, is in fact the B -immune property of the A -dominated set. A key step we take is to generalize the above argument to show that for the α -dense A -dominated set \mathcal{S}^A (which exists assuming that $\text{OPT}_B(\Pi) = 1 - \alpha < 1$), it holds that:

$$\Pr_{\langle A^{(1)}, B \rangle} [\mathcal{S}^A] \geq \frac{\alpha}{\text{val}(\Pi)}, \quad (2)$$

where $\text{val}(\Pi')$ is the expected outcome of Π' . Namely, in $(A^{(1)}, B)$ the probability of hitting the set \mathcal{S}^A of 1-transcripts is larger by a factor of at least $\frac{1}{\text{val}(\Pi)}$ than the probability of hitting this set in the original protocol Π . Again, it is easy to see that the above holds if A controls the root of the tree and has a 1-transcript as a direct descendant; see Figure 1 for a concrete example. The proof of the general case can be found in Section 3.

Consider now the protocol $(A^{(1)}, B)$. In this protocol, the probability of hitting the set \mathcal{S}^A is at least $\frac{\alpha}{\text{val}(\Pi)}$, and clearly the set \mathcal{S}^A remains B -immune. Hence, we can apply Equation (2) again, to deduce that

$$\Pr_{\langle A^{(2)}, B \rangle} [\mathcal{S}^A] = \Pr_{\langle (A^{(1)}), B \rangle} [\mathcal{S}^A] \geq \frac{\Pr_{\langle A^{(1)}, B \rangle} [\mathcal{S}^A]}{\text{val}(A^{(1)}, B)} \geq \frac{\alpha}{\text{val}(\Pi) \cdot \text{val}(A^{(1)}, B)}. \quad (3)$$

Continuing it for κ iterations yields that

$$\text{val}(A^{(\kappa)}, B) \geq \Pr_{\langle A^{(\kappa)}, B \rangle} [\mathcal{S}^A] \geq \frac{\alpha}{\prod_{i=0}^{\kappa-1} \text{val}(A^{(i)}, B)}. \quad (4)$$

So, modulo some cheating,⁹ it seems that we are in good shape. Taking, for example, $\kappa = \log(\frac{1}{\alpha}) / \log(\frac{1}{0.9})$, Equation (4) yields that $\text{val}(A^{(\kappa)}, B) > 0.9$. Namely, if we assume that \mathcal{A} has an advantage over \mathcal{B} , then by recursively applying the biased-continuation attack for A enough

⁷They show that the same holds for the analogous attackers carrying out the biased-continuation attack towards zero.

⁸A key step in the analysis of Haitner and Omri [11] is to consider the “all-cheating protocol” $(A^{(1,1)}, B^{(1,1)})$, where $A^{(1,1)}$ plays against $B^{(1,1)}$ and they both carry out the biased-continuation attack trying to bias the outcome towards one. Since, and this is easy to verify, the expected outcome of $(A^{(1,1)}, B^{(1,1)})$ is one, using symmetry one can show that the expected outcome of either $(A^{(1,1)}, B)$ or $(A, B^{(1,1)})$ is at least $\frac{1}{\sqrt{2}}$, yielding a bias of $\frac{1}{\sqrt{2}} - \frac{1}{2}$. As mentioned in [11], symmetry cannot be used to prove a bias larger than $\frac{1}{\sqrt{2}} - \frac{1}{2}$.

⁹The actual argument is somewhat more complicated than the one given above. To ensure the above argument holds we need to consider measures over the 1-transcripts (and not sets). In addition, while (the measure variant of) Equation (3) is correct, deriving it from Equation (2) takes some additional steps.




Figure 1: Coin-flipping protocol Π . The label of an internal node (i.e., partial transcript) denotes the name of the party controlling it (i.e., the party that sends the next message given this partial transcript), and that of a leaf (i.e., full transcript) denotes its value — the parties’ common output once reaching this leaf. Finally, the label on an edge leaving a node u to node u' denotes the probability that a random execution of Π visits u' once in u .

Note that $\text{OPT}_A(\Pi) = 1$ and $\text{OPT}_B(\Pi) = 1 - \alpha_1$. The A -dominated set \mathcal{S}^A in this case consists of the single 1-leaf to the left of the root. The conditional protocol Π' is the protocol rooted in the node to the right of the root (of Π), and the B' -dominated set \mathcal{S}^B consists of the single 0-leaf to the left of the root of Π' .

times, we arbitrarily bias the expected output of the protocol towards one. Unfortunately, if this advantage (i.e., $\alpha = (1 - \text{OPT}_B(\Pi))$) is very small, which is the case in typical examples, the number of recursions required might be linear in the protocol depth (or even larger). Given the recursive nature of the above attack, the running time of the described attacker is *exponential*. To overcome this obstacle, we consider not only the dominated set, but additional sets that are “close to” being dominated. Informally, we can say that a 1-transcript belongs to the A -dominated set if it can be generated by an execution of (\mathcal{A}, B) . In other words, the probability, over B ’s coins, that a transcript generated by a random execution of (\mathcal{A}, B) belongs to the A -dominated set is one. We define a set of 1-transcripts that does not belong to the A -dominated set to be “close to” A -dominated if there is an (unbounded) attacker $\hat{\mathcal{A}}$, such that the probability, over B ’s coins, that a transcript generated by a random execution of $(\hat{\mathcal{A}}, B)$ belongs to the set is close to one. These sets are formally defined via the notion of conditional protocols, discussed next.

Conditional Protocols. Let $\Pi = (A, B)$ be a coin-flipping protocol in which there exists an A -dominated set \mathcal{S}^A of density $\alpha > 0$. Consider the “conditional” protocol $\Pi' = (A', B')$, resulting from conditioning on not hitting the set \mathcal{S}_A . Namely, the message distribution of Π' is that induced by a random execution of Π that does not generate transcripts in \mathcal{S}_A . See Figure 1 for a concrete example. We note that the protocol Π' might not be efficiently computable (even if Π is), but this does not bother us, since we only use it as a thought experiment.

We have effectively removed all the 1-transcripts dominated by A (the set \mathcal{S}^A must contain all such transcripts; otherwise $\text{OPT}_B(\Pi)$ would be smaller than $1 - \alpha$). Thus, the expected outcome of (A', B') is zero. Therefore, $\text{OPT}_{B'}(\Pi') = 1$ and $\text{OPT}_{A'}(\Pi') = 1 - \beta < 1$. It follows from this crucial observation that there exists a B' -dominated \mathcal{S}^B of density β , over the 0-transcripts of Π' . Applying a similar argument to that used for Equation (4) yields that for large enough κ , the

biased-continuation attacker $B^{(\kappa)}$, playing the role of B' , succeeds in biasing the outcome of Π' toward zero, where κ is proportional to $\log(\frac{1}{\beta})$. Moreover, if α is small, the above yields that $B^{(\kappa)}$ does almost equally well in the original protocol Π . If β is also small, we can now consider the conditional protocol Π'' , obtained by conditioning Π' on not hitting the B' -dominated set, and so on.

By iterating the above process enough times, the A -dominated sets cover all the 1-transcripts, and the B -dominated sets cover all the 0-transcripts.¹⁰ Assume that in the above iterated process, the density of the A -dominated sets is the first to go beyond $\varepsilon > 0$. It can be shown — and this a key technical contribution of this paper — that it is almost as good as if the density of the *initial* set \mathcal{S}_A was ε .¹¹ We conclude that for any $\varepsilon > 0$, there exists a constant κ such that $\text{val}(A^{(\kappa)}, B) > 1 - \varepsilon$.¹²

1.3.2 Using the Transcript Inverter

We have seen above that for any constant ε , by recursively applying the biased-continuation attack for constantly many times, we get an attack that biases the outcome of the protocol by $\frac{1}{2} - \varepsilon$. The next thing is to implement the above attack *efficiently*, under the assumption that one-way functions do not exist. Given a partial transcript u of protocol Π , we wish to return a uniformly chosen full transcript of Π that is consistent with u and the common outcome it induces is one. Biased continuation can be reduced to the task of finding *honest continuation*: returning a uniformly chosen full transcript of Π that is consistent with u . Assuming honest continuation can be found for the protocol, biased-continuation can also be found by calling the honest continuation many times, until a transcript whose output is one is obtained. The latter can be done efficiently, as long as the value of the partial transcript u — the expected outcome of the protocol conditioned on u , is not too low. (If it is too low, too much time might pass before a full transcript leading to one is obtained.) Ignoring this low value problem, and noting that honest continuation of a protocol can be reduced to inverting the protocol's transcript function, all we need to do to implement $A^{(i)}$ is to invert the transcript functions of the protocols $(A, B), (A^{(1)}, B), \dots, (A^{(i-1)}, B)$. Furthermore, noting that the attackers $A^{(1)}, \dots, A^{(i-1)}$ are *stateless*, it suffices to have the ability to invert *only* the transcript function of (A, B) .

So attacking a coin-flipping protocol Π boils down to inverting the transcript function f_Π of Π , and making sure we are not doing that on low value transcripts. Assuming one-way functions do not exist, there exists an efficient inverter Inv for f_Π that is guaranteed to work well when invoked on random outputs of f_Π (i.e., when f_Π is invoked on the uniform distribution; nothing is guaranteed for distributions far from uniform). By the above discussion, algorithm Inv implies an efficient approximation of $A^{(i)}$, as long as the partial transcripts attacked by $A^{(i)}$ are neither *low-value* nor *unbalanced* (by low-value transcript we mean that the expected outcome of the protocol conditioned on the transcript is low; by unbalanced transcript we mean that its density with respect to $(A^{(i)}, B)$ is not too far from its density with respect to (A, B)). Whereas the authors of [11] proved that the queries of $A^{(1)}$ obey the two conditions with sufficiently high probability, we were unable to prove this (and believe it is untrue) for the queries of $A^{(i)}$, for $i > 1$. Thus, we simply cannot argue

¹⁰When considering measures and not sets, as done in the actual proof, this covering property is not trivial.

¹¹More accurately, let $\tilde{\mathcal{S}}^A$ be the union of these 1-transcript sets and let $\tilde{\alpha}$ be the density of $\tilde{\mathcal{S}}^A$ in Π . Then $\text{val}(A^{(\kappa)}, B) \geq \Pr_{(A^{(\kappa)}, B)} \left[\tilde{\mathcal{S}}^A \right] \geq \frac{\tilde{\alpha}}{\prod_{i=0}^{\kappa-1} \text{val}(A^{(i)}, B)}$.

¹²The assumption that the density of the A -dominated sets is the first to go beyond $\varepsilon > 0$ is independent of the assumption that \mathcal{A} wins in the zero-sum game $(\mathcal{A}, \mathcal{B})$. Specifically, the fact that $A^{(\kappa)}$ succeeds in biasing the protocol does not guarantee that \mathcal{A} is the winner of $(\mathcal{A}, \mathcal{B})$.

that $A^{(i)}$ has an efficient approximation, assuming one-way functions do not exist. Fortunately, we managed to prove the above for the “pruned” variant of $A^{(i)}$, defined below.

Unbalanced and low-value transcripts. Before defining our final attacker, we relate the problem of unbalanced transcripts to that of low-value transcripts. We say that a (partial) transcript u is γ -*unbalanced* if the probability that u is visited with respect to a random execution of $(A^{(1)}, B)$ is at least γ times larger than with respect to a random execution of (A, B) . Furthermore, we say that a (partial) transcript u is δ -*small* if the expected outcome of (A, B) , conditioned on visiting u , is at most δ . We prove (a variant of) the following statement. For any $\delta > 0$ and $\gamma > 1$, there exists c that depends on δ , such that

$$\Pr_{\ell \leftarrow (A^{(1)}, B)} [\ell \text{ has a } \gamma\text{-unbalanced prefix but no } \delta\text{-small prefix}] \leq \frac{1}{\gamma^c}. \quad (5)$$

Namely, as long as $(A^{(1)}, B)$ does not visit low-value transcript, it is only at low risk to significantly deviate (in a multiplicative sense) from the distribution induced by (A, B) . Equation (5) naturally extends to recursive biased-continuation attacks. It also has an equivalent form for the attacker $B^{(1)}$, trying to bias the protocol towards zero, with respect to δ -high transcripts — the expected outcome of Π , conditioned on visiting the transcript, is at least $1 - \delta$.

The pruning attacker. At last we are ready to define our final attacker. To this end, for protocol $\Pi = (A, B)$ we define its δ -*pruned variant* $\Pi_\delta = (A_\delta, B_\delta)$, where $\delta \in (0, \frac{1}{2})$, as follows. As long as the execution does not visit a δ -low or δ -high transcript, the parties act as in Π . Once a δ -low transcript is visited, only the party B sends messages, and it does so according to the distribution induced by Π . If a δ -high transcript is visited (and has no δ -low prefix), only the party A sends messages, and again it does so according to the distribution induced by Π .

Since the transcript distribution induced by Π_δ is the same as of Π , protocol Π_δ is also a coin-flipping protocol. We also note that Π_δ can be implemented efficiently assuming one-way functions do not exist (simply use the inverter of Π 's transcript function to estimate the value of a given transcript). Finally, by Equation (5), $A_\delta^{(i)}$ (i.e., recursive biased-continuation attacks for Π_δ) can be efficiently implemented, since there are *no* low-value transcripts where A needs to send the next message. (Similarly, $B_\delta^{(i)}$ can be efficiently implemented since there are no high-value transcripts where B needs to send the next message.)

It follows that for any constant $\varepsilon > 0$, there exists constant κ such that either the expected outcome of $(A_\delta^{(\kappa)}, B_\delta)$ is at least $1 - \varepsilon$, or the expected outcome of $(A_\delta, B_\delta^{(\kappa)})$ is at most ε . Assume for concreteness that it is the former case. We define our pruning attacker $A^{(\kappa, \delta)}$ as follows. When playing against B , the attacker $A^{(\kappa, \delta)}$ acts like $A_\delta^{(\kappa)}$ would when playing against B_δ . Namely, the attacker pretends that it is in the δ -pruned protocol Π_δ . But once a low- or high-value transcript is reached, $A^{(\kappa, \delta)}$ acts *honestly* in the rest of the execution (like A would).

It follows that until a low- or high-value transcript has been reached for the first time, the distribution of $(A^{(\kappa, \delta)}, B)$ is the same as that of $(A_\delta^{(\kappa)}, B_\delta)$. Once a δ -low transcript is reached, the expected outcome of both $(A^{(\kappa, \delta)}, B)$ and $(A_\delta^{(\kappa)}, B_\delta)$ is δ , but when a δ -high transcript is reached, the expected outcome of $(A^{(\kappa, \delta)}, B)$ is $(1 - \delta)$ (since it plays like A would), where the expected outcome of $(A_\delta^{(\kappa)}, B_\delta)$ is at most one. All in all, the expected outcome of $(A^{(\kappa, \delta)}, B)$ is δ -close to that of $(A_\delta^{(\kappa)}, B_\delta)$, and thus the expected outcome of $(A^{(\kappa, \delta)}, B)$ is at least $1 - \varepsilon - \delta$. Since ε and

δ are arbitrary constants, we have established an efficient attacker to bias the outcome of Π by a value that is an arbitrary constant close to one.

1.4 Open Questions

Does the existence of any non-trivial coin-flipping protocol (i.e., bias $\frac{1}{2} - \frac{1}{\text{poly}(n)}$) imply the existence of one-way functions? This is the main question left open. Answering it would fully resolve the computational complexity of coin-flipping protocols.

Paper Organization

General notations and definitions used throughout the paper are given in Section 2. Our ideal attacker (which has access to a perfect sampler) to bias any coin-flipping protocol is presented and analyzed in Section 3, while in Section 4 we show how to modify the above attack to be useful when the ideal attacker is replaced with a one-way function inverter.

Acknowledgment

We are very grateful to Hemanta Maji, Yishay Mansour, Eran Omri and Alex Samorodnitsky for useful discussions.

2 Preliminaries

2.1 Notations

We use calligraphic letters to denote sets, uppercase for random variables and functions, lowercase for values, boldface for vectors, and sans-serif (e.g., A) for algorithms (i.e., Turing Machines). All logarithms considered here are in base two, where \circ denotes string concatenation. Let \mathbb{N} denote the set of natural numbers, where 0 is considered as a natural number, i.e., $\mathbb{N} = \{0, 1, 2, 3, \dots\}$. For $n \in \mathbb{N}$, let $(n) = \{0, \dots, n\}$ and if n is positive let $[n] = \{1, \dots, n\}$, where $[0] = \emptyset$. For $a \in \mathbb{R}$ and $b \geq 0$, let $[a \pm b]$ stand for the interval $[a - b, a + b]$, $(a \pm b)$ for $(a - b, a + b)$ etc. For a non-empty string $t \in \{0, 1\}^*$ and $i \in [|t|]$, let t_i be the i 'th bit of t , and for $i, j \in [|t|]$ such that $i < j$, let $t_{i, \dots, j} = t_i \circ t_{i+1} \circ \dots \circ t_j$. The empty string is denoted by λ , and for a non-empty string, let $t_{1, \dots, 0} = \lambda$. We let poly denote the set all polynomials and let PPTM denote a probabilistic algorithm that runs in *strictly* polynomial time. Given a PPTM algorithm A , we let $\mathsf{A}(u; r)$ be an execution of A on input u given randomness r . A function $\nu: \mathbb{N} \mapsto [0, 1]$ is *negligible*, denoted $\nu(n) = \text{neg}(n)$, if $\nu(n) < 1/p(n)$ for every $p \in \text{poly}$ and large enough n .

Given a random variable X , we write $x \leftarrow X$ to indicate that x is selected according to X . Similarly, given a finite set \mathcal{S} , we let $s \leftarrow \mathcal{S}$ denote that s is selected according to the uniform distribution on \mathcal{S} . We adopt the convention that when the same random variable occurs several times in an expression, all occurrences refer to a single sample. For example, $\Pr[f(X) = X]$ is defined to be the probability that when $x \leftarrow X$, we have $f(x) = x$. We write U_n to denote the random variable distributed uniformly over $\{0, 1\}^n$. The support of a distribution D over a finite set \mathcal{U} , denoted $\text{Supp}(D)$, is defined as $\{u \in \mathcal{U} : D(u) > 0\}$. The *statistical distance* of two distributions P and Q over a finite set \mathcal{U} , denoted as $\text{SD}(P, Q)$, is defined as $\max_{\mathcal{S} \subseteq \mathcal{U}} |P(\mathcal{S}) - Q(\mathcal{S})| = \frac{1}{2} \sum_{u \in \mathcal{U}} |P(u) - Q(u)|$.

A *measure* is a function $M: \Omega \mapsto [0, 1]$. The support of M over a set Ω , denoted $\text{Supp}(M)$, is defined as $\{\omega \in \Omega: M(\omega) > 0\}$. A measure M over Ω is the *zero measure* if $\text{Supp}(M) = \emptyset$.

2.2 Two-Party Protocols

The following discussion is restricted to no-input (possibly randomized), two-party protocols, where each message consists of a *single* bit. We do not assume, however, that the parties play in turns (i.e., the same party might send two consecutive messages), but only that the protocol's transcript uniquely determines which party is playing next (i.e., the protocol is well defined). In an m -round protocol, the parties interact for exactly m rounds. The tuple of the messages sent so far in any partial execution of a protocol is called the (*communication*) *transcript* of this execution.

We write that a protocol Π is equal to (A, B) , when A and B are the interactive Turing Machines that control the left- and right-hand party respectively, of the interaction according to Π . For a party C interacting according to Π , let \bar{C}_Π be the other party in Π , where if Π is clear from the context, we simply write \bar{C} .

If A and B are deterministic, then $\text{trans}(A, B)$ denotes the uniquely defined transcript of the protocol (A, B) . If A and B are randomized, we let ρ_A and ρ_B be the (maximal) number of random bits used by A and B respectively. For $r_A \in \{0, 1\}^{\rho_A}$, $A(\cdot; r_A)$ stands for the variant of A when r_A are set as its random coins, and $A(u; r_A)$ is the message sent by $A(\cdot; r_A)$ when given a partial transcript u , for which the party A sends the next message. The above notations naturally extend for the party B as well. The transcript of the protocol $(A(\cdot; r_A), B(\cdot; r_B))$ is denoted by $\text{trans}(A(\cdot; r_A), B(\cdot; r_B))$. For a (partial) transcript u of a protocol $\Pi = (A, B)$, let $\text{Consis}_\Pi(u)$ be the distribution of choosing $(r_A, r_B) \leftarrow \{0, 1\}^{\rho_A} \times \{0, 1\}^{\rho_B}$ conditioned on $\text{trans}(A(\cdot; r_A), B(\cdot; r_B))_{1, \dots, |u|} = u$.

2.2.1 Binary Trees

Definition 2.1 (binary trees). For $m \in \mathbb{N}$, let \mathcal{T}^m be the complete directed binary tree of height m . We naturally identify the vertices of \mathcal{T}^m with binary strings: the root is denoted by the empty string λ , and the left- and right-hand children of a non-leaf node u are denoted by $u0$ and $u1$ respectively.

- Let $\mathcal{V}(\mathcal{T}^m)$, $\mathcal{E}(\mathcal{T}^m)$, $\text{root}(\mathcal{T}^m)$ and $\mathcal{L}(\mathcal{T}^m)$ denote the vertices, edges, root and leaves of \mathcal{T}^m respectively.
- For $u \in \mathcal{V}(\mathcal{T}^m) \setminus \mathcal{L}(\mathcal{T}^m)$, let \mathcal{T}_u^m be the sub-tree of \mathcal{T}^m rooted at u .
- For $u \in \mathcal{V}(\mathcal{T}^m)$, let $\text{desc}_m(u)$ [resp., $\overline{\text{desc}}_m(u)$] be the descendants of u in \mathcal{T}^m including u [resp., excluding u], and for $\mathcal{U} \subseteq \mathcal{V}(\mathcal{T}^m)$ let $\text{desc}_m(\mathcal{U}) = \bigcup_{u \in \mathcal{U}} \text{desc}_m(u)$ and $\overline{\text{desc}}_m(\mathcal{U}) = \bigcup_{u \in \mathcal{U}} \overline{\text{desc}}_m(u)$.
- The frontier of a set $\mathcal{U} \subseteq \mathcal{V}(\mathcal{T}^m)$, denoted by $\text{frnt}(\mathcal{U})$, is defined as $\mathcal{U} \setminus \overline{\text{desc}}_m(\mathcal{U})$.

When m is clear from the context, it is typically omitted from the above notation.

2.2.2 Protocol Trees

We naturally identify a (possibly partial) transcript of an m -round, single-bit message protocol with a rooted path in \mathcal{T}^m . That is, the transcript $t \in \{0, 1\}^m$ is identified with the path $\lambda, t_1, t_{1,2}, \dots, t$.

Definition 2.2 (tree representation of a protocol). *We make use of the following definitions with respect to an m -round protocol $\Pi = (A, B)$, and $C \in \{A, B\}$.*

- *Let $\text{round}(\Pi) = m$, let $\mathcal{T}(\Pi) = \mathcal{T}^m$, and for $X \in \{\mathcal{V}, \mathcal{E}, \text{root}, \mathcal{L}\}$ let $X(\Pi) = X(\mathcal{T}(\Pi))$.*
- *The edge distribution induced by a protocol Π is the function $e_\Pi: \mathcal{E}(\Pi) \mapsto [0, 1]$, defined as $e_\Pi(u, v)$ being the probability that the transcript of a random execution of Π visits v , conditioned that it visits u .*
- *For $u \in \mathcal{V}(\Pi)$, let $v_\Pi(u) = e_\Pi(\lambda, u_1) \cdot e_\Pi(u_1, u_{1,2}) \dots \cdot e_\Pi(u_{1,\dots,|u|-1}, u)$, and let the leaf distribution induced by Π be the distribution $\langle \Pi \rangle$ over $\mathcal{L}(\Pi)$, defined by $\langle \Pi \rangle(u) = v_\Pi(u)$.*
- *The party that sends the next message on transcript u is said to control u , and we denote this party by $\text{cntrl}_\Pi(u)$. We call $\text{cntrl}_\Pi: \mathcal{V}(\Pi) \mapsto \{A, B\}$ the control scheme of Π . Let $\text{Ctrl}_\Pi^C = \{u \in \mathcal{V}(\Pi): \text{cntrl}_\Pi(u) = C\}$.*

For $\mathcal{S} \subseteq \mathcal{V}(\Pi)$, let $\Pr_{\langle \Pi \rangle}[\mathcal{S}]$ be abbreviation for $\Pr_{\ell \leftarrow \langle \Pi \rangle}[\ell \in \mathcal{S}]$. Note that every function $e: \mathcal{E}(\mathcal{T}^m) \mapsto [0, 1]$ with $e(u, u0) + e(u, u1) = 1$ for every $u \in \mathcal{V}(\mathcal{T}^m) \setminus \mathcal{L}(\mathcal{T}^m)$ with $v(u) > 0$, along with a control scheme (active in each node), defines a two party, m -round, single-bit message protocol (the resulting protocol might be inefficient). The analysis in Section 3 naturally gives rise to functions over binary trees that do not correspond to any two-party execution. We identify the “protocols” induced by such functions by the special symbol \perp . We let $E_{\langle \perp \rangle}[f] = 0$, for any real-value function f .

The view of a protocol as an edge-distribution function allows us to consider protocols induced by sub-trees of $\mathcal{T}(\Pi)$.

Definition 2.3 (sub-protocols). *Let Π be a protocol and let $u \in \mathcal{V}(\Pi)$. Let $(\Pi)_u$ denote the protocol induced by the function e_Π on the sub-tree of $\mathcal{T}(\Pi)$ rooted at u , if such a protocol exists, and let $(\Pi)_u = \perp$ otherwise.*

Namely, the protocol $(\Pi)_u$ is the protocol Π conditioned on u being the transcript of the first $|u|$ rounds. When convenient, we remove the parentheses from notation, and simply write Π_u . Two sub-protocols of interest are Π_0 and Π_1 , induced by e_Π and the trees rooted at the left- and right-hand descendants of $\text{root}(\mathcal{T})$. For a measure $M: \mathcal{L}(\Pi) \mapsto [0, 1]$ and $u \in \mathcal{V}(\Pi)$, let $(M)_u: \mathcal{L}(\Pi_u) \mapsto [0, 1]$ be the restricted measure induced by M on the sub-protocol Π_u . Namely, for any $\ell \in \mathcal{L}(\Pi_u)$, $(M)_u(\ell) = M(\ell)$.

2.2.3 Tree Value

Definition 2.4 (tree value). *Let Π a two-party protocol that at the end of any of its executions, the parties output the same real value. Let $\chi_\Pi: \mathcal{L}(\Pi) \mapsto \mathbb{R}$ be the common output function of Π — $\chi_\Pi(\ell)$ is the common output of the parties in an execution ending in ℓ .¹³ Let $\text{val}(\Pi) = E_{\langle \Pi \rangle}[\chi_\Pi]$, and for $x \in \mathbb{R}$ let $\mathcal{L}_x(\Pi) = \{\ell \in \mathcal{L}(\Pi): \chi_\Pi(\ell) = x\}$.*

¹³Conditioned that an execution of the protocol generates a transcript ℓ , the parties’ coins are in a product distribution. Hence, if the parties always have the same output, then the protocol’s output is indeed a (deterministic) function of its transcript.

Throughout this paper we restrict ourselves to protocols whose common output is either one or zero, i.e., the image of χ_Π is the set $\{0, 1\}$. The following immediate fact states that the expected value of a measure, whose support is a subset of the 1-leaves of some protocol, is always smaller than the value of that protocol.

Fact 2.5. *Let Π be a protocol and let M be a measure over $\mathcal{L}_1(\Pi)$. Then $E_{\langle \Pi \rangle} [M] \leq \text{val}(\Pi)$.*

2.2.4 Protocol with Common Inputs

We sometimes would like to apply the above terminology to a protocol $\Pi = (A, B)$ whose parties get a common security parameter 1^n . This is formally done by considering the protocol $\Pi_n = (A_n, B_n)$, where C_n is the algorithm derived by “hardwiring” 1^n into the code of C .

2.3 Coin-Flipping Protocols

In a coin-flipping protocol two parties interact and in the end have a common output bit. Ideally, this bit should be random and no cheating party should be able to bias its outcome to either direction (if the other party remains honest). For interactive, probabilistic algorithms A and B , and $x \in \{0, 1\}^*$, let $\text{out}(A, B)(x)$ denote the parties’ output, on common input x .

Definition 2.6 ((strong) coin-flipping). *A PPT protocol (A, B) is a δ -bias coin-flipping protocol if the following holds.*

Correctness: $\Pr[\text{out}(A, B)(1^n) = (0, 0)] = \Pr[\text{out}(A, B)(1^n) = (1, 1)] = \frac{1}{2}$.

Security: $\Pr[\text{out}(A^*, B)(1^n) = (*, c)], \Pr[\text{out}(A, B^*)(1^n) = (c, *)] \leq \frac{1}{2} + \delta(n)$, for any PPTM’s A^* and B^* , bit $c \in \{0, 1\}$ and large enough n .

Sometimes, e.g., if the parties have (a priori known) opposite preferences, an even weaker definition of coin-flipping protocols is of interest.

Definition 2.7 (weak coin-flipping). *A PPT protocol (A, B) is a weak δ -bias coin-flipping protocol if the following holds.*

Correctness: Same as in Definition 2.6.

Security: There exist bits $c_A \neq c_B \in \{0, 1\}$ such that

$$\Pr[\text{out}(A^*, B)(1^n) = c_A], \Pr[\text{out}(A, B^*)(1^n) = c_B] \leq \frac{1}{2} + \delta(n)$$

for any PPTM’s A^* and B^* , and large enough n .

Remark 2.8. *Our result still holds when replacing the value $\frac{1}{2}$ in the correctness requirement above with any constant in $(0, 1)$. It also holds for protocols in which, with some small probability, the parties are not in agreement regarding the protocol’s outcome, or even might output values that are not bits.*

In the rest of the paper we restrict our attention to m -round single-bit message coin-flipping protocols, where $m = m(n)$ is a function of the protocol’s security parameter. Given such a protocol $\Pi = (A, B)$, we assume that its common output (i.e., the coin) is efficiently computable from a (full) transcript of the protocol. (It is easy to see that these assumptions are without loss of generality.)

2.4 One-Way Functions and Distributional One-Way Functions

A one-way function (OWF) is an efficiently computable function whose inverse cannot be computed on average by any PPTM.

Definition 2.9. A polynomial-time computable function $f: \{0, 1\}^n \mapsto \{0, 1\}^{\ell(n)}$ is one-way if

$$\Pr_{x \leftarrow \{0, 1\}^n; y = f(x)} [A(1^n, y) \in f^{-1}(y)] = \text{neg}(n)$$

for any PPTM A .

A seemingly weaker definition is that of a distributional OWF. Such a function is easy to compute, but it is hard to compute uniformly random preimages of random images.

Definition 2.10. A polynomial-time computable $f: \{0, 1\}^n \mapsto \{0, 1\}^{\ell(n)}$ is distributional one-way, if $\exists p \in \text{poly}$ such that

$$\text{SD}((x, f(x))_{x \leftarrow \{0, 1\}^n}, (A(f(x)), f(x))_{x \leftarrow \{0, 1\}^n}) \geq \frac{1}{p(n)}$$

for any PPTM A and large enough n .

Clearly, any one-way function is also a distributional one-way function. While the other implication is not necessarily always true, Impagliazzo and Luby [15] showed that the existence of distributional one-way functions implies that of (standard) one-way functions. In particular, the authors of [15] proved that if one-way functions do not exist, then any efficiently computable function has an inverter of the following form.

Definition 2.11 (ξ -inverter). An algorithm Inv is an ξ -inverter of $f: \mathcal{D} \mapsto \mathcal{R}$ if the following holds.

$$\Pr_{x \leftarrow \mathcal{D}; y = f(x)} [\text{SD}((x')_{x' \leftarrow f^{-1}(y)}, (\text{Inv}(y))) > \xi] \leq \xi.$$

Lemma 2.12 ([15, Lemma 1]). Assume one-way functions do not exist. Then for any polynomial-time computable function $f: \{0, 1\}^n \mapsto \{0, 1\}^{\ell(n)}$ and $p \in \text{poly}$, there exists a PPTM algorithm Inv such that the following holds for infinitely many n 's. On security parameter 1^n , algorithm Inv is a $1/p(n)$ -inverter of f_n (i.e., f is restricted to $\{0, 1\}^n$).

Impagliazzo and Luby [15] only gave a proof sketch for the above lemma. The full proof can be found in [14, Theorem 4.2.2].

Remark 2.13 (Definition of inverter). In their original definition, Impagliazzo and Luby [15] defined a ξ -inverter as an algorithm Inv for which it holds that

$$\text{SD}((x, f(x))_{x \leftarrow \{0, 1\}^n}, (\text{Inv}(f(x)), f(x))_{x \leftarrow \{0, 1\}^n}) < \xi.$$

They also proved Lemma 2.12 with respect to this definition. By taking, for example, $\xi' = \xi^2$ and applying their proof with ξ' , it is easy to see how our version of Lemma 2.12 follows with respect to the above definition of a ξ -inverter.

Note that nothing is guaranteed when invoking a good inverter (i.e., a γ -inverter for some small γ) on an arbitrary distribution. Yet the following lemma yields that if the distribution in consideration is “not too different” from the output distribution of f , then such good inverters are useful.

Lemma 2.14. *Let f and g be two randomized functions over the same domain $\mathcal{D} \cup \{\perp\}$ such that $f(\perp) \equiv g(\perp)$, and let $\{D_i\}_{i \in [k]}$ be a set of distributions over $\mathcal{D} \cup \{\perp\}$ such that for some $a \geq 0$ it holds that $\mathbb{E}_{d \leftarrow D_i}[\text{SD}(f(d), g(d))] \leq a$ for every $i \in [k]$. Let \mathbf{A} be a k -query oracle-aided algorithm that only makes queries in \mathcal{D} . Let $Q = (Q_1, \dots, Q_k)$ be the random variable of the queries of \mathbf{A}^f in such a random execution, setting $Q_i = \perp$ if \mathbf{A} makes less than i queries.*

Assume that $\Pr_{(q_1, \dots, q_k) \leftarrow Q} [\exists i \in [k]: q_i \neq \perp \wedge Q_i(q_i) > \lambda \cdot D_i(q_i)] \leq b$ for some $\lambda, b \geq 0$. Then $\text{SD}(\mathbf{A}^f, \mathbf{A}^g) \leq b + ka\lambda$.

To prove Lemma 2.14, we use the following proposition.

Proposition 2.15. *For every two distributions P and Q over a set \mathcal{D} , there exists a distribution $R_{P,Q}$ over $\mathcal{D} \times \mathcal{D}$, such that the following hold:*

1. $(R_{P,Q})_1 \equiv P$ and $(R_{P,Q})_2 \equiv Q$, where $(R_{P,Q})_b$ is the projection of $R_{P,Q}$ into its b 'th coordinate.
2. $\Pr_{(x_1, x_2) \leftarrow R_{P,Q}} [x_1 \neq x_2] = \text{SD}(P, Q)$.

Proof. For every $x \in \mathcal{D}$, let $M(x) = \min\{P(x), Q(x)\}$, let $M_P(x) = P(x) - M(x)$ and $M_Q(x) = Q(x) - M(x)$. The distribution $R_{P,Q}$ is defined by the following procedure. With probability $\mu = \sum_{x \in \mathcal{D}} M(x)$, sample an element x according to M (i.e., x is returned with probability $\frac{M(x)}{\mu}$), and return (x, x) ; otherwise return (x_P, x_Q) where x_P is sampled according to M_P and x_Q is sampled according to M_Q . It is clear that $\Pr_{(x_1, x_2) \leftarrow R_{P,Q}} [x_1 \neq x_2] = \text{SD}(P, Q)$. It also holds that

$$\begin{aligned} (R_{P,Q})_1(x) &= \mu \cdot \frac{M(x)}{\mu} + (1 - \mu) \cdot \frac{M_P(x)}{\mu_P} \\ &= M(x) + M_P(x) \\ &= P(x), \end{aligned}$$

where $\mu_P := \sum_{x \in \mathcal{D}} M_P = (1 - \mu)$. Namely, $(R_{P,Q})_1 \equiv P$. The proof that $(R_{P,Q})_2 \equiv Q$ is analogous. \square

Proof of Lemma 2.14. Using Proposition 2.15 and standard argument, it holds that $\text{SD}(\mathbf{A}^f, \mathbf{A}^g)$ is at most the probability that the following experiment aborts.

Experiment 2.16.

1. Start emulating a random execution of \mathbf{A} .
2. Do until \mathbf{A} halts:
 - (a) Let q be the next query of \mathbf{A} .
 - (b) Sample $(a_1, a_2) \leftarrow R_{f(q), g(q)}$.
 - (c) If $a_1 = a_2$, give a_1 to \mathbf{A} as the oracle answer.
Otherwise, abort.

By setting $\mathcal{S}_i = \{q : q \in \text{Supp}(Q_i) \wedge Q_i(q) \leq \lambda \cdot D_i(q)\}$ for $i \in [k]$ and recalling that by assumption $f(\perp) \equiv g(\perp)$ (thus, when sampling $(a_1, a_2) \leftarrow R_{f(\perp), g(\perp)}$, a_1 always equals a_2), we conclude that

$$\begin{aligned}
\text{SD}(A^f, A^g) &\leq \Pr_{(q_1, \dots, q_k) \leftarrow Q} [\exists i \in [k] : q_i \notin \mathcal{S}_i \cup \{\perp\}] \\
&\quad + \Pr_{(q_1, \dots, q_k) \leftarrow Q} [\exists i \in [k] : a_1 \neq a_2 \text{ where } (a_1, a_2) \leftarrow R_{f(q_i), g(q_i)} \wedge q_i \in \mathcal{S}_i] \\
&\leq b + \sum_{i \in [k]} \sum_{q \in \mathcal{S}_i} Q_i(q) \cdot \Pr[a_1 \neq a_2 \text{ where } (a_1, a_2) \leftarrow R_{f(q), g(q)}] \\
&\leq b + \sum_{i \in [k]} \sum_{q \in \mathcal{S}_i} Q_i(q) \cdot \text{SD}(f(q), g(q)) \\
&\leq b + \sum_{i \in [k]} \sum_{q \in \text{Supp}(D_i)} \lambda \cdot D_i(q) \cdot \text{SD}(f(q), g(q)) \\
&\leq b + \lambda \sum_{i \in [k]} \mathbb{E}_{q \leftarrow D_i} [\text{SD}(f(q), g(q))] \\
&\leq b + ka\lambda,
\end{aligned}$$

where the third inequality follows from Proposition 2.15 and the fourth from the definition of the sets $\{\mathcal{S}_i\}_{i \in [k]}$. \square

2.5 Two Inequalities

We make use of following technical lemmas, whose proofs are given in Appendix A.

Lemma 2.17. *Let $x, y \in [0, 1]$ and $a_1, \dots, a_k, b_1, \dots, b_k \in (0, 1]$. Then for any $p_0, p_1 \geq 0$ with $p_0 + p_1 = 1$, it holds that*

$$p_0 \cdot \frac{x^{k+1}}{\prod_{i=1}^k a_i} + p_1 \cdot \frac{y^{k+1}}{\prod_{i=1}^k b_i} \geq \frac{(p_0 x + p_1 y)^{k+1}}{\prod_{i=1}^k (p_0 a_i + p_1 b_i)}.$$

Lemma 2.18. *For every $\delta \in (0, \frac{1}{2}]$, there exists $\alpha = \alpha(\delta) \in (0, 1]$ such that*

$$\lambda \cdot a_1^{1+\alpha} \cdot (2 - a_1 \cdot x) + a_2^{1+\alpha} \cdot (2 - a_2 \cdot x) \leq (1 + \lambda) \cdot (2 - x),$$

for every $x \geq \delta$ and $\lambda, y \geq 0$ with $\lambda y \leq 1$, for $a_1 = 1 + y$ and $a_2 = 1 - \lambda y$.

3 The Biased-Continuation Attack

In this section we describe an attack to bias any coin-flipping protocol. The described attack, however, might be impossible to implement efficiently (even when assuming one-way functions do not exist). Specifically, we assume access to an ideal sampling algorithm to sample a *uniform* preimage of *any* output of the functions under consideration. Our actual attack, the subject of Section 4, tries to mimic the behavior of this attack while being efficiently implemented (assuming one-way functions do not exist).

The following discussion is restricted to (coin-flipping) protocols whose parties always output the same bit as their common output, and this bit is determined by the protocol's transcript.

In all protocols considered in this section, the messages are bits. In addition, the protocols under consideration have no inputs (neither private nor common), and in particular no security parameter is involved.¹⁴ Recall that \perp stands for a canonical invalid/undefined protocol, and that $E_{\langle \perp \rangle}[f] = 0$, for any real value function f . (We refer the reader to Section 2 for a discussion of the conventions and assumptions used above.) Although the focus of this paper is coin-flipping protocols, all the results in this section hold true for any two-party protocol meeting the above assumptions. Specifically, we do not assume that an honest execution of the protocol produces a uniformly random bit, nor do we assume that the parties executing the protocol can be implemented by a polynomial time probabilistic Turing machine. For this reason we omit the term “coin-flipping” in this section.

Throughout the paper we prove statements with respect to attackers that, when playing the role of the left-hand party of the protocol (i.e., A), are trying to bias the common output of the protocol towards one, and, when playing the role of the right-hand party of the protocol (i.e., B), are trying to bias the common output of the protocol towards zero. All statements have analogues ones with respect to the opposite attack goals.

Let $\Pi = (A, B)$ be a protocol. The *recursive biased-continuation attack* described below recursively applies the *biased-continuation attack* introduced by Haitner and Omri [11].¹⁵ The biased-continuation attacker $A_{\Pi}^{(1)}$ – playing the role of A – works as follows: in each of A ’s turns, $A_{\Pi}^{(1)}$ picks a random continuation of Π , whose output it induces is equal to one, and plays the current turn accordingly. The i ’th biased-continuation attacker $A_{\Pi}^{(i)}$, formally described below, uses the same strategy but the random continuation taken is of the protocol $(A_{\Pi}^{(i-1)}, B)$.

Moving to the formal discussion, for a protocol $\Pi = (A, B)$, we let BiasedCont_{Π} be the following algorithm.

Definition 3.1 (BiasedCont_{Π}).

Input: $u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi)$ and a bit $b \in \{0, 1\}$

Operation:

1. Choose $\ell \leftarrow \langle \Pi \rangle$ conditioned that
 - (a) $\ell \in \text{desc}(u)$, and
 - (b) $\chi_{\Pi}(\ell) = b$.¹⁶
2. Return $\ell_{|u|+1}$.

Let $A_{\Pi}^{(0)} \equiv A$, and for integer $i > 0$ define:

Algorithm 3.2 ($A_{\Pi}^{(i)}$).

Input: transcript $u \in \{0, 1\}^*$.

Operation:

1. If $u \in \mathcal{L}(\Pi)$, output $\chi_{\Pi}(u)$ and halt.

¹⁴In Section 4, we make use of these input-less protocols by “hardwiring” the security parameter of the protocols under consideration.

¹⁵Called the “random continuation attack” in [11].

¹⁶If no such ℓ exists, the algorithm returns an arbitrary leaf in $\text{desc}(u)$.

2. Set $\text{msg} = \text{BiasedCont}_{(A_{\Pi}^{(i-1)}, B)}(u, 1)$.
3. Send msg to B .
4. If $u' = u \circ \text{msg} \in \mathcal{L}(\Pi)$, output $\chi_{\Pi}(u')$.¹⁷

The attacker $B_{\Pi}^{(i)}$ attacking towards zero is analogously defined (specifically, the call $\text{BiasedCont}_{(A_{\Pi}^{(i-1)}, B)}(u, 1)$ in Algorithm 3.2 is changed to $\text{BiasedCont}_{(A, B_{\Pi}^{(i-1)})}(u, 0)$).¹⁸

It is relatively easy to show that the more recursions $A_{\Pi}^{(i)}$ and $B_{\Pi}^{(i)}$ do, the closer their success probability is to that of an all-powerful attacker, who can either bias the outcome to zero or to one. The important point of the following theorem is that, for any $\varepsilon > 0$, there exists a *global* constant $\kappa = \kappa(\varepsilon)$ (i.e., independent of the underlying protocol), for which either $A_{\Pi}^{(\kappa)}$ or $B_{\Pi}^{(\kappa)}$ succeeds in its attack with probability at least $1 - \varepsilon$. This becomes crucial when trying to efficiently implement these adversaries (see Section 4), as each recursion call might induce a polynomial blowup in the running time of the adversary. Since κ is constant (for a constant ε), the recursive attacker is still efficient.

Theorem 3.3 (main theorem, ideal version). *For every $\varepsilon \in (0, \frac{1}{2}]$ there exists an integer $\kappa = \kappa(\varepsilon) \geq 0$ such that for every protocol $\Pi = (A, B)$, either $\text{val}(A_{\Pi}^{(\kappa)}, B) > 1 - \varepsilon$ or $\text{val}(A, B_{\Pi}^{(\kappa)}) < \varepsilon$.*

The rest of this section is devoted to proving the above theorem.

In what follows, we typically omit the subscript Π from the notation of the above attackers. Towards proving Theorem 3.3 we show a strong (and somewhat surprising) connection between recursive biased-continuation attacks on a given protocol and the optimal valid attack on this protocol. The latter is the best (unbounded) attack on this protocol, which sends only valid messages (ones that could have been sent by the honest party). Towards this goal we define sequences of measures over the leaves (i.e., transcripts) of the protocol, connect these measures to the optimal attack, and then lower bound the success of the recursive biased-continuation attacks using these measures.

In the following we first observe some basic properties of the recursive biased-continuation attack. Next, we define the optimal valid attack, define a simple measure with respect to this attack, and analyze, as a warm-up, the success of recursive biased-continuation attacks on this measure. After arguing why considering the latter measure does not suffice, we define a sequence of measures, and then state, in Section 3.6, a property of this sequence that yields Theorem 3.3 as a corollary. The main body of this section deals with proving Section 3.6,

3.1 Basic Observations About $A^{(i)}$

We make two basic observations regarding the recursive biased-continuation attack. The first gives expression to the edge distribution this attack induces. The second is that this attack is stateless.

¹⁷For the mere purpose of biasing B 's output, there is no need for $A^{(i)}$ to output anything. Yet doing so helps us to simplify our recursion definitions (specifically, we use the fact that in $(A^{(i)}, B)$ the parties always have the same output).

¹⁸The subscript Π is added to the notation (i.e., $A_{\Pi}^{(i)}$), since the biased-continuation attack for A depends not only on the definition of the party A , but also on the definition of B , the other party in the protocol.

We'll use these observations in the following sections; however, the reader might want to skip their straightforward proofs for now.

Recall that at each internal node in its control, $A^{(1)}$ picks a random continuation to one. We can also describe $A^{(1)}$'s behavior as follows: after seeing a transcript u , $A^{(1)}$ biases the probability of sending, e.g., 0 to B : it does so proportionally to the ratio between the chance of having output one among all honest executions of the protocol that are consistent with the transcript $u \circ 0$, and the same chance but with respect to the transcript u . The behavior of $A^{(i)}$ is analogous where $A^{(i-1)}$ replaces the role of A in the above discussion. Formally, we have the following claim.

Claim 3.4. *Let $\Pi = (A, B)$ be a protocol and let $A^{(j)}$ be according to Algorithm 3.2. Then*

$$e_{(A^{(i)}, B)}(u, ub) = e_{\Pi}(u, ub) \cdot \frac{\prod_{j=0}^{i-1} \text{val}((A^{(j)}, B)_{ub})}{\prod_{j=0}^{i-1} \text{val}((A^{(j)}, B)_u)}, \quad 19$$

for any $i \in \mathbb{N}$, A -controlled $u \in \mathcal{V}(\Pi)$ and $b \in \{0, 1\}$.

This claim is a straightforward generalization of the proof of [11, Lemma 12]. However, for completeness and to give an example of our notations, a full proof is given below.

Proof. The proof is by induction on i . For $i = 0$, recall that $A^{(0)} \equiv A$, and hence $e_{(A^{(0)}, B)}(u, ub) = e_{\Pi}(u, ub)$, as required.

Assume the claim holds for $i - 1$, and we want to compute $e_{(A^{(i)}, B)}(u, ub)$. The definition of Algorithm 3.2 yields that for any positive $i \in \mathbb{N}$, it holds that

$$\begin{aligned} e_{(A^{(i)}, B)}(u, ub) &= \Pr_{\ell \leftarrow \langle A^{(i-1)}, B \rangle} \left[\ell_{|u|+1} = b \mid \ell \in \text{desc}(u) \wedge \chi_{(A^{(i-1)}, B)}(\ell) = 1 \right] \quad 20 \\ &= \frac{\Pr_{\ell \leftarrow \langle A^{(i-1)}, B \rangle} \left[\ell_{|u|+1} = b \wedge \chi_{(A^{(i-1)}, B)}(\ell) = 1 \mid \ell \in \text{desc}(u) \right]}{\Pr_{\ell \leftarrow \langle A^{(i-1)}, B \rangle} \left[\chi_{(A^{(i-1)}, B)}(\ell) = 1 \mid \ell \in \text{desc}(u) \right]} \\ &= e_{(A^{(i-1)}, B)}(u, ub) \cdot \frac{\text{val}((A^{(i-1)}, B)_{ub})}{\text{val}((A^{(i-1)}, B)_u)}, \end{aligned} \quad (6)$$

where the last equality is by a simple chain rule, i.e., since

$$\begin{aligned} e_{(A^{(i-1)}, B)}(u, ub) &= \Pr_{\ell \leftarrow \langle A^{(i-1)}, B \rangle} \left[\ell_{|u|+1} = b \mid \ell \in \text{desc}(u) \right], \text{ and} \\ \text{val}((A^{(i-1)}, B)_{ub}) &= \Pr_{\ell \leftarrow \langle A^{(i-1)}, B \rangle} \left[\chi_{(A^{(i-1)}, B)}(\ell) = 1 \mid \ell \in \text{desc}(u) \wedge \ell_{|u|+1} = b \right]. \end{aligned}$$

The proof is concluded by plugging the induction hypothesis into Equation (6). \square

The following observation enables us to use induction when analyzing the power of $A^{(i)}$.

Proposition 3.5. *For every protocol $\Pi = (A_{\Pi}, B_{\Pi})$, $i \in \mathbb{N}$ and $b \in \{0, 1\}$, it holds that $(A_{\Pi}^{(i)}, B)_{b_i}$ and $(A_{\Pi_b}^{(i)}, B_{\Pi_b})$ are the same protocol, where $\Pi_b = (A_{\Pi_b}, B_{\Pi_b})$.*

¹⁹Recall that for a protocol Π and a partial transcript u , we let $e_{\Pi}(u, ub)$ stand for the probability that the party controlling u sends b as the next message, conditioning that u is the transcript of the execution thus far.

²⁰Recall that for a protocol Π , we let $\langle \Pi \rangle$ stand for the leaf distribution of Π .

Proof. Immediately follows from $A_{\Pi}^{(i)}$ being stateless. \square

Remark 3.6. Note that the party B_{Π_b} , defined by the subprotocol Π_b (specifically, by the edge distribution of the subtree $\mathcal{T}(\Pi_b)$), might not have an efficient implementation, even if B does have one. For the sake of the arguments we make in this section, however, it matters only that B_{Π_b} is well defined.

3.2 Optimal Valid Attacks

When considering the optimal attackers for a given protocol, we restrict ourselves to valid attackers. Informally, we can say that, on each of its turns, a valid attacker sends a message from the set of possible replies that the honest party might choose given the transcript so far.

Definition 3.7 (optimal valid attacker). Let $\Pi = (A, B)$ be a protocol. A deterministic algorithm A' playing the role of A in Π is in \mathcal{A}^* , if $v_{\Pi}(u) = 0 \implies v_{(A', B)}(u) = 0$ for any $u \in \mathcal{V}(\Pi)$. The class \mathcal{B}^* is analogously defined. Let $\text{OPT}_A(\Pi) = \max_{A' \in \mathcal{A}^*} \{\text{val}(A', B)\}$ and $\text{OPT}_B(\Pi) = \max_{B' \in \mathcal{B}^*} \{1 - \text{val}(A, B')\}$.

The following proposition is immediate.

Proposition 3.8. Let $\Pi = (A, B)$ be a protocol and let $u \in \mathcal{V}(\Pi)$. Then,

$$\text{OPT}_A(\Pi_u) = \begin{cases} \chi_{\Pi}(u) & u \in \mathcal{L}(\Pi); \\ \max \{\text{OPT}_A(\Pi_{ub}) : e_{\Pi}(u, ub) > 0\}, & u \notin \mathcal{L}(\Pi) \text{ and } u \text{ is controlled by } A; \\ e_{\Pi}(u, u0) \cdot \text{OPT}_A(\Pi_{u0}) + e_{\Pi}(u, u1) \cdot \text{OPT}_A(\Pi_{u1}), & u \notin \mathcal{L}(\Pi) \text{ and } u \text{ is controlled by } B, \end{cases}$$

and the analog conditions hold for $\text{OPT}_B(\Pi_u)$.²¹

The following holds true for any (bit value) protocol.

Proposition 3.9. Let $\Pi = (A, B)$ be a protocol with $\text{val}(\Pi) \in [0, 1]$. Then either $\text{OPT}_A(\Pi)$ or $\text{OPT}_B(\Pi)$ (but not both) is equal to 1.

The somewhat surprising part is that *only* one party has a valid winning strategy. Assume for simplicity that $\text{OPT}_A(\Pi) = 1$. Since A might accidentally mimic the optimal winning valid attacker, it follows that for any valid strategy B' for B there is a positive probability over the random choices of the honest A that the outcome is *not* zero. Namely, it holds that $\text{OPT}_B(\Pi) < 1$. The formal proof follows a straightforward induction on the protocol's round complexity.

Proof of Proposition 3.9. The proof is by induction on the round complexity of Π . Assume that $\text{round}(\Pi) = 0$ and let ℓ be the only node in $\mathcal{T}(\Pi)$. If $\chi_{\Pi}(\ell) = 1$, the proof follows since $\text{OPT}_A(\Pi) = 1$ and $\text{OPT}_B(\Pi) = 0$. In the complementary case, i.e., $\chi_{\Pi}(\ell) = 0$, the proof follows since $\text{OPT}_A(\Pi) = 0$ and $\text{OPT}_B(\Pi) = 1$.

Assume that the lemma holds for m -round protocols and that $\text{round}(\Pi) = m + 1$. If $e_{\Pi}(\lambda, b) = 1$ ²² for some $b \in \{0, 1\}$, since Π is a protocol, it holds that $e_{\Pi}(\lambda, 1 - b) = 0$. Hence, by

²¹Recall that for a (possible partial) transcript u , Π_u is the protocol Π , conditioned that $u_1, \dots, u_{|u|}$ were the first $|u|$ messages.

²²Recall that λ is the string representation of the root of $\mathcal{T}(\Pi)$.

Proposition 3.8 it holds that $\text{OPT}_A(\Pi) = \text{OPT}_A(\Pi_b)$ and $\text{OPT}_B(\Pi) = \text{OPT}_B(\Pi_b)$, regardless of the party controlling $\text{root}(\Pi)$. The proof follows from the induction hypothesis.

If $e_\Pi(\lambda, b) \notin \{0, 1\}$ for both $b \in \{0, 1\}$, the proof splits according to the following complementary cases:

$\text{OPT}_B(\Pi_0) < 1$ **and** $\text{OPT}_B(\Pi_1) < 1$. The induction hypothesis yields that $\text{OPT}_A(\Pi_0) = 1$ and $\text{OPT}_A(\Pi_1) = 1$. Proposition 3.8 now yields that $\text{OPT}_B(\Pi) < 1$ and $\text{OPT}_A(\Pi) = 1$, regardless of the party controlling $\text{root}(\Pi)$.

$\text{OPT}_B(\Pi_0) = 1$ **and** $\text{OPT}_B(\Pi_1) = 1$. The induction hypothesis yields that $\text{OPT}_A(\Pi_0) < 1$ and $\text{OPT}_A(\Pi_1) < 1$. Proposition 3.8 now yields that $\text{OPT}_B(\Pi) = 1$ and $\text{OPT}_A(\Pi) < 1$, regardless of the party controlling $\text{root}(\Pi)$.

$\text{OPT}_B(\Pi_0) = 1$ **and** $\text{OPT}_B(\Pi_1) < 1$. The induction hypothesis yields that $\text{OPT}_A(\Pi_0) < 1$ and $\text{OPT}_A(\Pi_1) = 1$. If A controls $\text{root}(\Pi)$, Proposition 3.8 yields that $\text{OPT}_A(\Pi) = 1$ and $\text{OPT}_B(\Pi) < 1$. If B controls $\text{root}(\Pi)$, Proposition 3.8 yields that $\text{OPT}_A(\Pi) < 1$ and $\text{OPT}_B(\Pi) = 1$. Hence, the proof follows.

$\text{OPT}_B(\Pi_0) < 1$ and $\text{OPT}_B(\Pi_1) = 1$. The proof follows arguments similar to the previous case. \square

In the next sections we show the connection between the optimal valid attack and recursive biased-continuation attacks, by connecting them both to a specific measure over the protocol's leaves, called here the ‘‘dominated measure’’ of a protocol.

3.3 Dominated Measures

Consider the following measure over the protocol's leaves.

Definition 3.10 (dominated measures). *The A-dominated measure of protocol $\Pi = (A, B)$, denoted M_Π^A , is a measure over $\mathcal{L}(\Pi)$ defined as $M_\Pi^A(\ell) = \chi_\Pi(\ell)$ if $\text{round}(\Pi) = 0$, and otherwise recursively defined by:*

$$M_\Pi^A(\ell) = \begin{cases} 0, & e_\Pi(\lambda, \ell_1) = 0;^{23} \\ M_{\Pi_{\ell_1}}^A(\ell_2, \dots, |\ell|), & e_\Pi(\lambda, \ell_1) = 1; \\ M_{\Pi_{\ell_1}}^A(\ell_2, \dots, |\ell|), & e_\Pi(\lambda, \ell_1) \notin \{0, 1\} \wedge (A \text{ controls } \text{root}(\Pi) \vee \text{Smaller}_\Pi(\ell_1)); \\ \frac{\mathbb{E}_{\langle \Pi_{1-\ell_1} \rangle} [M_{\Pi_{1-\ell_1}}^A]}{\mathbb{E}_{\langle \Pi_{\ell_1} \rangle} [M_{\Pi_{\ell_1}}^A]} \cdot M_{\Pi_{\ell_1}}^A(\ell_2, \dots, |\ell|), & \text{otherwise,} \end{cases}$$

where $\text{Smaller}_\Pi(\ell_1) = 1$ if $\mathbb{E}_{\langle \Pi_{\ell_1} \rangle} [M_{\Pi_{\ell_1}}^A] \leq \mathbb{E}_{\langle \Pi_{1-\ell_1} \rangle} [M_{\Pi_{1-\ell_1}}^A]$. Finally, we let M_\perp^A be the zero measure.

The B-dominated measure of protocol Π , denoted M_Π^B , is analogously defined, except that $M_\Pi^B(\ell) = 1 - \chi_\Pi(\ell)$ if $\text{round}(\Pi) = 0$.

The following key observation justifies the name of the above measures.

²³Recall that for transcript ℓ , ℓ_1 stands for the first messages sent in ℓ .

Lemma 3.11. *Let $\Pi = (A, B)$ be a protocol and let M_{Π}^A be its A-dominated measure. Then $\text{OPT}_B(\Pi) = 1 - \mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A]$.*

In particular, since $\text{OPT}_A(\Pi) = 1$ iff $\text{OPT}_B(\Pi) < 1$ (Proposition 3.8), it holds that $\text{OPT}_A(\Pi) = 1$ iff $\mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A] > 0$.

The proof of Lemma 3.11 is given below. For the intuitive explanation, note that if A controls the root, the expected value of the A-dominated measure is the weighted average of the measures of the subprotocols Π_0 and Π_1 (according to the edge distributions). However, if B controls the root, the expected value is that of the lowest measure of the same subprotocols. Hence, in both cases the A-dominated measure “captures” the behavior of the optimal adversary for B.

Example 3.12. *Before continuing with the formal proof, we believe the reader might find the following concrete example useful. Let $\Pi = (A, B)$ be the protocol described in Figure 2a and assume for the sake of this example that $\alpha_0 < \alpha_1$. The A-dominated measures of Π and its subprotocols are given in Figure 2b.*

We would like to highlight some points regarding the calculations of the A-dominated measures. The first point we note is that $M_{\Pi_{011}}^A(011) = 1$ but $M_{\Pi_{01}}^A(011) = 0$. Namely, the A-dominated measure of the subprotocol Π_{011} assigns the leaf represented by the string 011 with the value 1, while the A-dominated measure of the subprotocol Π_{01} (for which Π_{011} is a subprotocol) assigns the same leaf with the value 0. This follows since $\mathbb{E}_{\langle \Pi_{010} \rangle} [M_{\Pi_{010}}^A] = 0$ and $\mathbb{E}_{\langle \Pi_{011} \rangle} [M_{\Pi_{011}}^A] = 1$, which yield that $\text{Smaller}_{\Pi_{01}}(1) = 0$ (recall that $\text{Smaller}_{\Pi'}(b) = 0$ iff the expected value of the A-dominated measure of Π'_b is larger than that of the A-dominated measure of Π'_{1-b}). Hence, Definition 3.10 with respect to Π_{01} now yields that

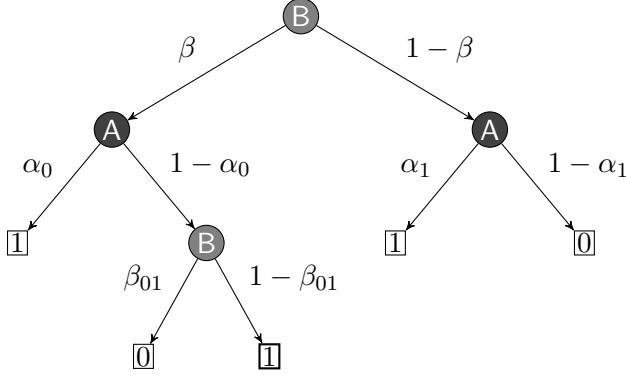
$$\begin{aligned} M_{\Pi_{01}}^A(011) &= \frac{\mathbb{E}_{\langle \Pi_{010} \rangle} [M_{\Pi_{010}}^A]}{\mathbb{E}_{\langle \Pi_{011} \rangle} [M_{\Pi_{011}}^A]} \cdot M_{\Pi_{011}}^A(011) \\ &= \frac{0}{1} \cdot 1 = 0. \end{aligned}$$

The second point we note is that $M_{\Pi_1}^A(10) = 1$ but $M_{\Pi}^A(10) = \frac{\alpha_0}{\alpha_1}$ (recall that we assumed that $\alpha_0 < \alpha_1$, so $\frac{\alpha_0}{\alpha_1} < 1$). This follows similar arguments to the previous point; it holds that $\mathbb{E}_{\langle \Pi_0 \rangle} [M_{\Pi_0}^A] = \alpha_0$ and $\mathbb{E}_{\langle \Pi_1 \rangle} [M_{\Pi_1}^A] = \alpha_1$, which yield that $\text{Smaller}_{\Pi}(1) = 0$ (since $\alpha_0 < \alpha_1$). Definition 3.10 with respect to Π now yields that

$$\begin{aligned} M_{\Pi}^A(10) &= \frac{\mathbb{E}_{\langle \Pi_0 \rangle} [M_{\Pi_0}^A]}{\mathbb{E}_{\langle \Pi_1 \rangle} [M_{\Pi_1}^A]} \cdot M_{\Pi_1}^A(10) \\ &= \frac{\alpha_0}{\alpha_1} \cdot 1 = \frac{\alpha_0}{\alpha_1}. \end{aligned}$$

The third and final point we note is the implication of Lemma 3.11 for this protocol. By the assumption that $\alpha_0 < \alpha_1$, it holds that $\text{OPT}_B(\Pi) = 1 - \alpha_0$. Independently, let us calculate the expected value of the A-dominated measure. Since $\text{Supp}(M_{\Pi}^A) = \{00, 01\}$, it holds that

$$\begin{aligned} \mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A] &= v_{\Pi}(00) \cdot M_{\Pi}^A(00) + v_{\Pi}(10) \cdot M_{\Pi}^A(10) \\ &= \beta \cdot \alpha_0 \cdot 1 + (1 - \beta) \cdot \alpha_1 \cdot \frac{\alpha_0}{\alpha_1} \\ &= \alpha_0. \end{aligned}$$



(a) Protocol $\Pi = (A, B)$. The label of an internal node denotes the name of the party controlling it, and that of a leaf denotes its value. The label on an edge leaving a node u to node u' denotes the probability that a random execution of Π visits u' once in u . Finally, all nodes are represented as strings from the root of Π , even when considering subprotocols (e.g., the string representations of the leaf with the thick borders is 011).

measures	Leaves				
	00	010	011	10	11
$M_{\Pi_{00}}^A$	1				
$M_{\Pi_{010}}^A$		0			
$M_{\Pi_{011}}^A$			1		
$M_{\Pi_{01}}^A$		0	0		
$M_{\Pi_0}^A$	1	0	0		
$M_{\Pi_{10}}^A$				1	
$M_{\Pi_{11}}^A$					0
$M_{\Pi_1}^A$				1	0
M_{Π}^A	1	0	0	α_0/α_1	0

(b) Calculating the A-dominated measure of Π . The A-dominated measure of a subprotocol Π_u , is only defined over the leaves in the subtree $\mathcal{T}(\Pi_u)$.

Figure 2: An example of a (coin-flipping) protocol is given on the left, and an example of how to calculate its A-dominated measure is given on the right.

Hence, $E_{\langle \Pi \rangle} [M_{\Pi}^A] = 1 - \text{OPT}_B(\Pi)$.

Towards proving Lemma 3.11, we first note that the definition of M_{Π}^A ensures three important properties.

Proposition 3.13. *Let Π be a protocol with $e_{\Pi}(\lambda, b) \notin \{0, 1\}$ for both $b \in \{0, 1\}$. Then*

1. (A-maximal) A controls $\text{root}(\Pi) \implies (M_{\Pi}^A)_b \equiv M_{\Pi_b}^A$ for both $b \in \{0, 1\}$.²⁴
2. (B-minimal) B controls $\text{root}(\Pi) \implies (M_{\Pi}^A)_b \equiv \begin{cases} M_{\Pi_b}^A, & \text{Smaller}_{\Pi}(b) = 1; \\ \frac{E_{\langle \Pi_{1-b} \rangle} [M_{\Pi_{1-b}}^A]}{E_{\langle \Pi_b \rangle} [M_{\Pi_b}^A]} \cdot M_{\Pi_b}^A, & \text{otherwise.} \end{cases}$
3. (B-immune) B controls $\text{root}(\Pi) \implies E_{\langle \Pi_0 \rangle} [(M_{\Pi}^A)_0] = E_{\langle \Pi_1 \rangle} [(M_{\Pi}^A)_1]$.

Namely, if A controls $\text{root}(\Pi)$, the A-maximal property of M_{Π}^A (the A-dominated measure of Π) ensures that the restrictions of this measure to the subprotocols of Π are the A-dominated measures of these subprotocols. In the complementary case, i.e., B controls $\text{root}(\Pi)$, the B-minimal property of M_{Π}^A ensures that for at least one subprotocol of Π , the restriction of this measure to this subprotocol is equal to the A-dominated measure of the subprotocol. Moreover, the B-immune property of M_{Π}^A ensures that the expected values of the measures derived by restricting M_{Π}^A to the subprotocols of Π are equal (and hence, they are also equal to the expected value of M_{Π}^A).

²⁴Recall that for a measure $M: \mathcal{L}(\Pi) \mapsto [0, 1]$ and a bit b , $(M)_b$ is the measure induced by M when restricted to $\mathcal{L}(\Pi_b) \subseteq \mathcal{L}(\Pi)$.

Proof of Proposition 3.13. The proof of Items 1 and 2 immediately follows from Definition 3.10.

Towards proving Item 3, we will assume that \mathbf{B} controls $\text{root}(\Pi)$. If $\text{Smaller}_\Pi(0) = \text{Smaller}_\Pi(1) = 1$, the proof again follows immediately from Definition 3.10. In the complementary case, i.e., $\text{Smaller}_\Pi(b) = 0$ and $\text{Smaller}_\Pi(1-b) = 1$ for some $b \in \{0, 1\}$, it holds that

$$\begin{aligned} \mathbb{E}_{\langle \Pi_b \rangle} \left[\left(M_\Pi^A \right)_b \right] &= \mathbb{E}_{\langle \Pi_b \rangle} \left[\frac{\mathbb{E}_{\langle \Pi_{1-b} \rangle} \left[M_{\Pi_{1-b}}^A \right]}{\mathbb{E}_{\langle \Pi_b \rangle} \left[M_{\Pi_b}^A \right]} \cdot M_{\Pi_b}^A \right] \\ &= \frac{\mathbb{E}_{\langle \Pi_{1-b} \rangle} \left[M_{\Pi_{1-b}}^A \right]}{\mathbb{E}_{\langle \Pi_b \rangle} \left[M_{\Pi_b}^A \right]} \cdot \mathbb{E}_{\langle \Pi_b \rangle} \left[M_{\Pi_b}^A \right] \\ &= \mathbb{E}_{\langle \Pi_{1-b} \rangle} \left[M_{\Pi_{1-b}}^A \right] \\ &= \mathbb{E}_{\langle \Pi_{1-b} \rangle} \left[\left(M_\Pi^A \right)_{1-b} \right], \end{aligned}$$

where the first and last equalities follow the \mathbf{B} -minimal property of M_Π^A (Proposition 3.13(2)). \square

We are now ready to prove Lemma 3.11.

Proof of Lemma 3.11. The proof is by induction on the round complexity of Π .

Assume that $\text{round}(\Pi) = 0$ and let ℓ be the only node in $\mathcal{T}(\Pi)$. If $\chi_\Pi(\ell) = 1$, then by Definition 3.10 it holds that $M_\Pi^A(\ell) = 1$, implying that $\mathbb{E}_{\langle \Pi \rangle} \left[M_\Pi^A \right] = 1$. The proof follows since in this case, by Proposition 3.9, $\text{OPT}_\mathbf{B}(\Pi) = 0$. In the complementary case, i.e., $\chi(\ell) = 0$, by Definition 3.10 it holds that $M_\Pi^A(\ell) = 0$, implying that $\mathbb{E}_{\langle \Pi \rangle} \left[M_\Pi^A \right] = 0$. The proof follows since in this case, by Proposition 3.9, $\text{OPT}_\mathbf{B}(\Pi) = 1$.

Assume that the lemma holds for m -round protocols and that $\text{round}(\Pi) = m+1$. For $b \in \{0, 1\}$ let $\alpha_b := \mathbb{E}_{\langle \Pi_b \rangle} \left[M_{\Pi_b}^A \right]$. The induction hypothesis yields that $\text{OPT}_\mathbf{B}(\Pi_b) = 1 - \alpha_b$ for both $b \in \{0, 1\}$. If $e_\Pi(\lambda, b) = 1$ for some $b \in \{0, 1\}$ (which also means that $e_\Pi(\lambda, 1-b) = 0$), the proof follows since Proposition 3.8 yields that $\text{OPT}_\mathbf{B}(\Pi) = \text{OPT}_\mathbf{B}(\Pi_b) = 1 - \alpha_b$, where Definition 3.10 yields that $\mathbb{E}_{\langle \Pi \rangle} \left[M_\Pi^A \right] = \mathbb{E}_{\langle \Pi_b \rangle} \left[M_{\Pi_b}^A \right] = \alpha_b$.

Assume $e_\Pi(\lambda, b) \notin \{0, 1\}$ for both $b \in \{0, 1\}$ and let $p := e_\Pi(\lambda, 0)$. The proof splits according to who controls the root of Π .

A controls $\text{root}(\Pi)$. Definition 3.10 yields that

$$\begin{aligned} \mathbb{E}_{\langle \Pi \rangle} \left[M_\Pi^A \right] &= p \cdot \mathbb{E}_{\langle \Pi_0 \rangle} \left[\left(M_\Pi^A \right)_0 \right] + (1-p) \cdot \mathbb{E}_{\langle \Pi_1 \rangle} \left[\left(M_\Pi^A \right)_1 \right] \\ &= p \cdot \mathbb{E}_{\langle \Pi_0 \rangle} \left[M_{\Pi_0}^A \right] + (1-p) \cdot \mathbb{E}_{\langle \Pi_1 \rangle} \left[M_{\Pi_1}^A \right] \\ &= p \cdot \alpha_0 + (1-p) \cdot \alpha_1, \end{aligned}$$

where the second equality follows the \mathbf{A} -maximal property of $M_{\Pi_b}^A$ (Proposition 3.13(1)).

Using Proposition 3.8 we conclude that

$$\begin{aligned}
\text{OPT}_B(\Pi) &= p \cdot \text{OPT}_B(\Pi_0) + (1-p) \cdot \text{OPT}_B(\Pi_1) \\
&= p \cdot (1 - \alpha_0) + (1-p) \cdot (1 - \alpha_1) \\
&= 1 - (p \cdot \alpha_0 + (1-p) \cdot \alpha_1) \\
&= 1 - \mathbb{E}_{\langle \Pi \rangle} \left[M_{\Pi}^A \right].
\end{aligned}$$

B controls root(Π). We assume that $\alpha_0 \leq \alpha_1$ (the complementary case is analogous). Proposition 3.8 and the induction hypothesis yield that $\text{OPT}_B(A, B) = 1 - \alpha_0$. Hence, it is left to show that $\mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A] = \alpha_0$. The assumption that $\alpha_0 \leq \alpha_1$ yields that $\text{Smaller}_{\Pi}(0) = 1$. Thus, by the B-minimal property of M_{Π}^A (Proposition 3.13(2)), it holds that $(M_{\Pi}^A)_0 \equiv M_{\Pi_0}^A$. It follows that $\mathbb{E}_{\langle \Pi_0 \rangle} [(M_{\Pi}^A)_0] = \alpha_0$, and the B-immune property of M_{Π}^A (Proposition 3.13(3)) yields that $\mathbb{E}_{\langle \Pi_1 \rangle} [(M_{\Pi}^A)_1] = \alpha_0$. To conclude the proof, we compute

$$\begin{aligned}
\mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A] &= p \cdot \mathbb{E}_{\langle \Pi_0 \rangle} \left[(M_{\Pi}^A)_0 \right] + (1-p) \cdot \mathbb{E}_{\langle \Pi_1 \rangle} \left[(M_{\Pi}^A)_1 \right] \\
&= p \cdot \alpha_0 + (1-p) \cdot \alpha_0 \\
&= \alpha_0.
\end{aligned}$$

□

Lemma 3.11 connects the success of the optimal attack to the expected value of the dominated measure. In the next section we analyze the success of the recursive biased-continuation attack using this expected value. Unfortunately, this analysis does not seem to suffice for our goal. In Section 3.5 we generalize the dominated measure described above to a sequence of (alternating) dominated measures, where in Section 3.6 we use this new notion to prove that the recursive biased continuation is indeed a good attack.

3.4 Warmup — Proof Attempt Using a (Single) Dominated Measure

As mentioned above, the approach described in this section falls too short to serve our goals. Yet we describe it here as a detailed overview for the more complicated proof, given in following sections (with respect to a sequence of dominated measures). Specifically, we sketch a proof of the following lemma, which relates the performance of the recursive biased-continuation attacker playing the role of A, to the performance of the optimal (valid) attacker playing the role of B. The proof (see below) is via the A-dominated measure of Π defined above.²⁵

Lemma 3.14. *Let $\Pi = (A, B)$ be a protocol with $\text{val}(\Pi) > 0$, let $k \in \mathbb{N}$ and let $A^{(k)}$ be according to Algorithm 3.2. Then*

$$\text{val}(A^{(k)}, B) \geq \frac{1 - \text{OPT}_B(\Pi)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)}.$$

The proof of the above lemma is a direct implication of the next lemma.

²⁵The formal proof of Lemma 3.14 follows its stronger variant, Lemma 3.25, introduced in Section 3.6.

Lemma 3.15. *Let $\Pi = (A, B)$ be a protocol with $\text{val}(\Pi) > 0$, let $k \in \mathbb{N}$ and let $A^{(k)}$ be according to Algorithm 3.2. Then*

$$\mathbb{E}_{\langle A^{(k)}, B \rangle} [M_{\Pi}^A] \geq \frac{\mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A]}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)}.$$

Proof of Lemma 3.14. Immediately follows Lemmas 3.11 and 3.15 and Fact 2.5. \square

We begin by sketching the proof of the following lemma, which is a special case of Lemma 3.15. Later we explain how to generalize the proof below to derive Lemma 3.15.

Lemma 3.16. *Let $\Pi = (A, B)$ be a protocol with $\text{val}(\Pi) > 0$ and let $A^{(1)}$ be according to Algorithm 3.2. Then $\mathbb{E}_{\langle A^{(1)}, B \rangle} [M_{\Pi}^A] \geq \frac{\mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A]}{\text{val}(\Pi)}$.*

Proof sketch. The proof is by induction on the round complexity of Π . The base case (i.e., $\text{round}(\Pi) = 0$) is straightforward. Assume that the lemma holds for m -round protocols and that $\text{round}(\Pi) = m + 1$. For $b \in \{0, 1\}$ let $\alpha_b := \mathbb{E}_{\langle \Pi_b \rangle} [M_{\Pi_b}^A]$ and let $p := e_{\Pi}(\lambda, 0)$.

If $\text{root}(\Pi)$ is controlled by A , the A -maximal property of M_{Π}^A (Proposition 3.13(1)) yields that $\mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A] = p \cdot \alpha_0 + (1 - p) \cdot \alpha_1$. It holds that

$$\begin{aligned} \mathbb{E}_{\langle A^{(1)}, B \rangle} [M_{\Pi}^A] &= e_{\langle A^{(1)}, B \rangle}(\lambda, 0) \cdot \mathbb{E}_{\langle (A^{(1)}, B)_0 \rangle} \left[\left(M_{\Pi}^A \right)_0 \right] + e_{\langle A^{(1)}, B \rangle}(\lambda, 1) \cdot \mathbb{E}_{\langle (A^{(1)}, B)_1 \rangle} \left[\left(M_{\Pi}^A \right)_1 \right] \\ &= p \cdot \frac{\text{val}(\Pi_0)}{\text{val}(\Pi)} \cdot \mathbb{E}_{\langle (A^{(1)}, B)_0 \rangle} \left[\left(M_{\Pi}^A \right)_0 \right] + (1 - p) \cdot \frac{\text{val}(\Pi_1)}{\text{val}(\Pi)} \cdot \mathbb{E}_{\langle (A^{(1)}, B)_1 \rangle} \left[\left(M_{\Pi}^A \right)_1 \right], \end{aligned} \quad (7)$$

where the second equality follows Claim 3.4. Since $A^{(1)}$ is stateless (Proposition 3.5), we can write Equation (7) as

$$\mathbb{E}_{\langle A^{(1)}, B \rangle} [M_{\Pi}^A] = p \cdot \frac{\text{val}(\Pi_0)}{\text{val}(\Pi)} \cdot \mathbb{E}_{\langle A_{\Pi_0}^{(1)}, B_{\Pi_0} \rangle} \left[\left(M_{\Pi}^A \right)_0 \right] + (1 - p) \cdot \frac{\text{val}(\Pi_1)}{\text{val}(\Pi)} \cdot \mathbb{E}_{\langle A_{\Pi_1}^{(1)}, B_{\Pi_1} \rangle} \left[\left(M_{\Pi}^A \right)_1 \right]. \quad (8)$$

The A -maximal property of M_{Π}^A and Equation (8) yield that

$$\mathbb{E}_{\langle A^{(1)}, B \rangle} [M_{\Pi}^A] = p \cdot \frac{\text{val}(\Pi_0)}{\text{val}(\Pi)} \cdot \mathbb{E}_{\langle A_{\Pi_0}^{(1)}, B_{\Pi_0} \rangle} [M_{\Pi_0}^A] + (1 - p) \cdot \frac{\text{val}(\Pi_1)}{\text{val}(\Pi)} \cdot \mathbb{E}_{\langle A_{\Pi_1}^{(1)}, B_{\Pi_1} \rangle} [M_{\Pi_1}^A]. \quad (9)$$

Applying the induction hypothesis on the right-hand side of Equation (9) yields that

$$\begin{aligned} \mathbb{E}_{\langle A^{(1)}, B \rangle} [M_{\Pi}^A] &\geq p \cdot \frac{\text{val}(\Pi_0)}{\text{val}(\Pi)} \cdot \frac{\alpha_0}{\text{val}(\Pi_0)} + (1 - p) \cdot \frac{\text{val}(\Pi_1)}{\text{val}(\Pi)} \cdot \frac{\alpha_1}{\text{val}(\Pi_1)} \\ &= \frac{p \cdot \alpha_0 + (1 - p) \cdot \alpha_1}{\text{val}(\Pi)} \\ &= \frac{\mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A]}{\text{val}(\Pi)}, \end{aligned}$$

which concludes the proof for the case that A controls $\text{root}(\Pi)$.

If $\text{root}(\Pi)$ is controlled by \mathbf{B} , and assuming that $\alpha_0 \leq \alpha_1$ (the complementary case is analogous), it holds that $\text{Smaller}_\Pi(0) = 1$. Thus, by the \mathbf{B} -minimal property of $M_\Pi^{\mathbf{A}}$ (Proposition 3.13(2)), it holds that $(M_\Pi^{\mathbf{A}})_0 \equiv M_{\Pi_0}^{\mathbf{A}}$ and $(M_\Pi^{\mathbf{A}})_1 \equiv \frac{\alpha_0}{\alpha_1} M_{\Pi_1}^{\mathbf{A}}$. Hence, the \mathbf{B} -immune property of $M_\Pi^{\mathbf{A}}$ (Proposition 3.13(3)) yields that $\mathbb{E}_{\langle \Pi \rangle} [M_\Pi^{\mathbf{A}}] = \alpha_0$. In addition, since \mathbf{B} controls $\text{root}(\Pi)$, the distribution of the edges $(\lambda, 0)$ and $(\lambda, 1)$ has not changed. It holds that

$$\begin{aligned} \mathbb{E}_{\langle \mathbf{A}^{(1)}, \mathbf{B} \rangle} [M_\Pi^{\mathbf{A}}] &= p \cdot \mathbb{E}_{\langle (\mathbf{A}^{(1)}, \mathbf{B})_0 \rangle} \left[(M_\Pi^{\mathbf{A}})_0 \right] + (1-p) \cdot \mathbb{E}_{\langle (\mathbf{A}^{(1)}, \mathbf{B})_1 \rangle} \left[(M_\Pi^{\mathbf{A}})_1 \right] \\ &= p \cdot \mathbb{E}_{\langle \mathbf{A}_{\Pi_0}^{(1)}, \mathbf{B}_{\Pi_0} \rangle} \left[(M_\Pi^{\mathbf{A}})_0 \right] + (1-p) \cdot \mathbb{E}_{\langle \mathbf{A}_{\Pi_1}^{(1)}, \mathbf{B}_{\Pi_1} \rangle} \left[(M_\Pi^{\mathbf{A}})_1 \right] \\ &= p \cdot \mathbb{E}_{\langle \mathbf{A}_{\Pi_0}^{(1)}, \mathbf{B}_{\Pi_0} \rangle} \left[M_{\Pi_0}^{\mathbf{A}} \right] + (1-p) \cdot \mathbb{E}_{\langle \mathbf{A}_{\Pi_1}^{(1)}, \mathbf{B}_{\Pi_1} \rangle} \left[\frac{\alpha_0}{\alpha_1} M_{\Pi_1}^{\mathbf{A}} \right] \\ &= p \cdot \mathbb{E}_{\langle \mathbf{A}_{\Pi_0}^{(1)}, \mathbf{B}_{\Pi_0} \rangle} \left[M_{\Pi_0}^{\mathbf{A}} \right] + (1-p) \cdot \frac{\alpha_0}{\alpha_1} \cdot \mathbb{E}_{\langle \mathbf{A}_{\Pi_1}^{(1)}, \mathbf{B}_{\Pi_1} \rangle} \left[M_{\Pi_1}^{\mathbf{A}} \right], \end{aligned} \quad (10)$$

where the second equality follows since $\mathbf{A}^{(1)}$ is stateless (Proposition 3.5). Applying the induction hypothesis on the right-hand side of Equation (10) yields that

$$\begin{aligned} \mathbb{E}_{\langle \mathbf{A}^{(1)}, \mathbf{B} \rangle} [M_\Pi^{\mathbf{A}}] &\geq p \cdot \frac{\alpha_0}{\text{val}(\Pi_0)} + (1-p) \cdot \frac{\alpha_0}{\alpha_1} \cdot \frac{\alpha_1}{\text{val}(\Pi_1)} \\ &= \alpha_0 \left(\frac{p}{\text{val}(\Pi_0)} + \frac{1-p}{\text{val}(\Pi_1)} \right) \\ &\geq \frac{\mathbb{E}_{\langle \Pi \rangle} [M_\Pi^{\mathbf{A}}]}{\text{val}(\Pi)}, \end{aligned}$$

which concludes the proof for the case that \mathbf{A} controls $\text{root}(\Pi)$, and where the last equality holds since

$$\frac{p}{\text{val}(\Pi_0)} + \frac{1-p}{\text{val}(\Pi_1)} \geq \frac{1}{\text{val}(\Pi)}. \quad (11)$$

□

The proof of Lemma 3.15 follows similar arguments to those used above for proving Lemma 3.16.²⁶ Informally, we proved Lemma 3.16 by showing that $\mathbf{A}^{(1)}$ “assigns” more weight to the dominated measure than \mathbf{A} does. A natural step is to consider $\mathbf{A}^{(2)}$ and to see if it assigns more weight to the dominated measure than $\mathbf{A}^{(1)}$ does. It turns out that one can turn this intuitive argument into a formal proof, and prove Lemma 3.14 by repeating this procedure with respect to many recursive biased-continuation attacks.²⁷

The shortcoming of Lemma 3.14. Given a protocol $\Pi = (\mathbf{A}, \mathbf{B})$, we are interested in the minimal value of κ for which $\mathbf{A}^{(\kappa)}$ biases the value of the protocol towards one with probability of at least 0.9 (as a concrete example). Following Lemma 3.14, it suffices to find a value κ such that

$$\text{val}(\mathbf{A}^{(\kappa)}, \mathbf{B}) \geq \frac{1 - \text{OPT}_{\mathbf{B}}(\Pi)}{\prod_{i=0}^{\kappa-1} \text{val}(\mathbf{A}^{(i)}, \mathbf{B})} \geq 0.9. \quad (12)$$

²⁶The proof sketch given for Lemma 3.16 is almost a formal proof, lacking only consideration of the base case and the extreme cases in which $e_\Pi(\lambda, b) = 1$ for some $b \in \{0, 1\}$.

²⁷The main additional complication in the proof of Lemma 3.14 is that the simple argument used to derive Equation (11) is replaced with the more general argument, described in Lemma 2.17.

Using worst case analysis, it suffices to find κ such that $(1 - \text{OPT}_{\mathbf{B}}(\Pi))/(0.9)^\kappa \geq 0.9$, where the latter dictates that

$$\kappa \geq \frac{\log\left(\frac{1}{1 - \text{OPT}_{\mathbf{B}}(\Pi)}\right)}{\log\left(\frac{1}{0.9}\right)}. \quad (13)$$

Recall that our ultimate goal is to implement an *efficient* attack on any coin-flipping protocol, under the mere assumption that one-way functions do not exist. Specifically, we would like to do so by giving an efficient version of the recursive biased-continuation attack. At the very least, due to the recursive nature of the attack, this requires the protocols $(A^{(1)}, B), \dots, (A^{(\kappa-1)}, B)$ be efficient in comparison to the basic protocol. The latter efficiency restriction together with the recursive definition of $A^{(\kappa)}$ dictates that κ (the number of recursion calls) be constant.

Unfortunately, Equation (13) reveals that if $\text{OPT}_{\mathbf{B}}(\Pi) \in 1 - o(1)$, we need to take $\kappa \in \omega(1)$, yielding an inefficient attack.

3.5 Back to the Proof — Sequence of Alternating Dominated Measures

Let $\Pi = (A, B)$ be a protocol and let M be a measure over the leaves of Π . Consider the variant of Π whose parties act identically to the parties in Π , but with the following tweak: when the execution reaches a leaf ℓ , the protocol restarts with probability $M(\ell)$. Namely, a random execution of the resulting (possibly inefficient) protocol is distributed like a random execution of Π , conditioned on not “hitting” the measure M .²⁸ The above is formally captured by the definition below.

3.5.1 Conditional Protocols

Definition 3.17 (conditional protocols). *Let Π be an m -message protocol and let M be a measure over $\mathcal{L}(\Pi)$ with $\mathbb{E}_{\langle \Pi \rangle}[M] < 1$. The m -message, M -conditional protocol of Π , denoted $\Pi|_{\neg M}$, is defined by the color function $\chi_{(\Pi|_{\neg M})} \equiv \chi_{\Pi}$, and the edge distribution function $e_{(\Pi|_{\neg M})}$ is defined by*

$$e_{(\Pi|_{\neg M})}(u, ub) = \begin{cases} 0, & \mathbb{E}_{\langle \Pi_u \rangle}[M] = 1;^{29} \\ e_{\Pi}(u, ub) \cdot \frac{1 - \mathbb{E}_{\langle \Pi_{ub} \rangle}[M]}{1 - \mathbb{E}_{\langle \Pi_u \rangle}[M]}, & \text{otherwise.} \end{cases},$$

for every $u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi)$ and $b \in \{0, 1\}$. The controlling scheme of the protocol $\Pi|_{\neg M}$ is the same as in Π .

If $\mathbb{E}_{\langle \Pi \rangle}[M] = 1$ or $\Pi = \perp$, we set $\Pi|_{\neg M} = \perp$.

The next proposition shows that the M -conditional protocol is indeed a protocol. It also shows a relation between the leaf distribution of the M -conditional protocol and the original protocol. Using this relation we conclude that the set of possible transcripts of the M -conditional protocol is a subset the original protocol’s possible transcripts and that if M assigns a value of 1 to some transcript, then this transcript is inaccessible by the M -conditional protocol.

²⁸For concreteness, one might like to consider the case where M is a set.

²⁹Note that this case does not affect the resulting protocol, and is defined only to simplify future discussion.

Proposition 3.18. *Let Π be a protocol and let M be a measure over $\mathcal{L}(\Pi)$ with $\mathbb{E}_{(\Pi)}[M] < 1$. Then*

1. $\forall u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi): \quad v_{(\Pi|\neg M)}(u) > 0 \implies e_{(\Pi|\neg M)}(u, u0) + e_{(\Pi|\neg M)}(u, u1) = 1;$
2. $\forall \ell \in \mathcal{L}(\Pi): \quad v_{(\Pi|\neg M)}(\ell) = v_{\Pi}(\ell) \cdot \frac{1 - M(\ell)}{1 - \mathbb{E}_{(\Pi)}[M]};$
3. $\forall \ell \in \mathcal{L}(\Pi): \quad v_{(\Pi|\neg M)}(\ell) > 0 \implies v_{\Pi}(\ell) > 0; \text{ and}$
4. $\forall \ell \in \mathcal{L}(\Pi): \quad M(\ell) = 1 \implies v_{(\Pi|\neg M)}(\ell) = 0.$

Proof. The first two items immediately follow from Definition 3.17. The last two items follow the second item. \square

In addition to the above properties, Definition 3.17 guarantees the following ‘‘locality’’ property of the M -conditional protocol.

Proposition 3.19. *Let Π be a protocol and let M be a measure over $\mathcal{L}(\Pi)$. Then $(\Pi|\neg M)_u = \Pi_u|\neg(M)_u$ for every $u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi)$.*

Proof. Immediately follows from Definition 3.17. \square

Proposition 3.19 helps us to apply induction on conditional protocols. Specifically, we use it to prove the following lemma, which relates the (dominated measure)-conditional protocol to the optimal (valid) attack.

Lemma 3.20. *Let $\Pi = (A, B)$ be a protocol with $\text{val}(\Pi) < 1$. Then $\text{OPT}_B(\Pi|\neg M_{\Pi}^A) = 1$.*

Proof. First, we note that Fact 2.5 yields that $\mathbb{E}_{(\Pi)}[M_{\Pi}^A] \leq \text{val}(\Pi) < 1$, and hence $\Pi|\neg M_{\Pi}^A \neq \perp$ (i.e., is a protocol). The rest of the proof is by induction on the round complexity of Π .

Assume that $\text{round}(\Pi) = 0$ and let ℓ be the only node in $\mathcal{T}(\Pi)$. Since it is assumed that $\text{val}(\Pi) < 1$, it must be the case that $\chi_{\Pi}(\ell) = 0$. The proof follows since $M_{\Pi}^A(\ell) = 0$, and thus $\Pi|\neg M_{\Pi}^A = \Pi$, and since $\text{OPT}_B(\Pi) = 1$.

Assume the lemma holds for m -round protocols and that $\text{round}(\Pi) = m + 1$. If $e_{\Pi}(\lambda, b) = 1$ for some $b \in \{0, 1\}$, Definition 3.10 yields that $(M_{\Pi}^A)_b = M_{\Pi_b}^A$. Moreover, Definition 3.17 yields that $e_{(\Pi|\neg M_{\Pi}^A)}(\lambda, b) = 1$. It holds that

$$\begin{aligned}
\text{OPT}_B(\Pi|\neg M_{\Pi}^A) &= \text{OPT}_B\left(\left(\Pi|\neg M_{\Pi}^A\right)_b\right) \\
&= \text{OPT}_B\left(\Pi_b|\neg\left(M_{\Pi}^A\right)_b\right) \\
&= \text{OPT}_B\left(\Pi_b|\neg M_{\Pi_b}^A\right) \\
&= 1,
\end{aligned} \tag{14}$$

where the first equality follows Proposition 3.8, the second follows from Proposition 3.19, and the last equality follows from the induction hypothesis.

In the complementary case, i.e., $e_{\Pi}(\lambda, b) \notin \{0, 1\}$ for both $b \in \{0, 1\}$, the proof splits according to who controls the root of Π .

A controls $\text{root}(\Pi)$. The assumption that $\text{val}(\Pi) < 1$ dictates that $\text{val}(\Pi_0) < 1$ or $\text{val}(\Pi_1) < 1$. Consider the following complimentary cases.

$\text{val}(\Pi_0), \text{val}(\Pi_1) < 1$: Proposition 3.8 yields that

$$\begin{aligned}
& \text{OPT}_{\mathbf{B}} \left(\Pi | \neg M_{\Pi}^{\mathbf{A}} \right) \\
&= e_{(\Pi | \neg M_{\Pi}^{\mathbf{A}})}(\lambda, 0) \cdot \text{OPT}_{\mathbf{B}} \left(\left(\Pi | \neg M_{\Pi}^{\mathbf{A}} \right)_0 \right) + e_{(\Pi | \neg M_{\Pi}^{\mathbf{A}})}(\lambda, 1) \cdot \text{OPT}_{\mathbf{B}} \left(\left(\Pi | \neg M_{\Pi}^{\mathbf{A}} \right)_1 \right) \\
&= e_{(\Pi | \neg M_{\Pi}^{\mathbf{A}})}(\lambda, 0) \cdot \text{OPT}_{\mathbf{B}} \left(\Pi_0 | \neg \left(M_{\Pi}^{\mathbf{A}} \right)_0 \right) + e_{(\Pi | \neg M_{\Pi}^{\mathbf{A}})}(\lambda, 1) \cdot \text{OPT}_{\mathbf{B}} \left(\Pi_1 | \neg \left(M_{\Pi}^{\mathbf{A}} \right)_1 \right) \\
&= e_{(\Pi | \neg M_{\Pi}^{\mathbf{A}})}(\lambda, 0) \cdot \text{OPT}_{\mathbf{B}} \left(\Pi_0 | \neg M_{\Pi_0}^{\mathbf{A}} \right) + e_{(\Pi | \neg M_{\Pi}^{\mathbf{A}})}(\lambda, 1) \cdot \text{OPT}_{\mathbf{B}} \left(\Pi_1 | \neg M_{\Pi_1}^{\mathbf{A}} \right) \\
&= 1,
\end{aligned}$$

where the first equality follows from Proposition 3.8, the second follows from Proposition 3.19, the third follows from by the **A**-maximal property of $M_{\Pi}^{\mathbf{A}}$ (Proposition 3.13(1)), and last equality follows from the induction hypothesis.

$\text{val}(\Pi_0) < 1, \text{val}(\Pi_1) = 1$: By Definition 3.17, it holds that

$$\begin{aligned}
e_{(\Pi | \neg M_{\Pi}^{\mathbf{A}})}(\lambda, 1) &= e_{\Pi}(\lambda, 1) \cdot \frac{1 - \mathbb{E}_{\langle \Pi_1 \rangle} \left[\left(M_{\Pi}^{\mathbf{A}} \right)_1 \right]}{1 - \mathbb{E}_{\langle \Pi \rangle} \left[M_{\Pi}^{\mathbf{A}} \right]} \\
&= e_{\Pi}(\lambda, 1) \cdot \frac{1 - \mathbb{E}_{\langle \Pi_1 \rangle} \left[M_{\Pi_1}^{\mathbf{A}} \right]}{1 - \mathbb{E}_{\langle \Pi \rangle} \left[M_{\Pi}^{\mathbf{A}} \right]} \\
&= 0,
\end{aligned}$$

where the second equality follows from the **A**-maximal property of $M_{\Pi}^{\mathbf{A}}$, and the last equality follows since $\text{val}(\Pi_1) = 1$, which yields that $\mathbb{E}_{\langle \Pi_1 \rangle} \left[M_{\Pi_1}^{\mathbf{A}} \right] = 1$. Since $\Pi | \neg M_{\Pi}^{\mathbf{A}}$ is a protocol (Proposition 3.18), it holds that $e_{(\Pi | \neg M_{\Pi}^{\mathbf{A}})}(\lambda, 0) = 1$. The proof now follows from Equation (14).

$\text{val}(\Pi_0) = 1, \text{val}(\Pi_1) < 1$: The proof in analogous to the previous case.

B controls $\text{root}(\Pi)$. Assume for simplicity that $\text{Smaller}_{\Pi}(0) = 1$, namely that $\mathbb{E}_{\langle \Pi_0 \rangle} \left[M_{\Pi_0}^{\mathbf{A}} \right] \leq \mathbb{E}_{\langle \Pi_1 \rangle} \left[M_{\Pi_1}^{\mathbf{A}} \right]$ (the other case is analogous). It must hold that $\text{val}(\Pi_0) < 1$ (otherwise, it holds that $\mathbb{E}_{\langle \Pi_0 \rangle} \left[M_{\Pi_0}^{\mathbf{A}} \right] = \mathbb{E}_{\langle \Pi_1 \rangle} \left[M_{\Pi_1}^{\mathbf{A}} \right] = 1$, which yields that $\text{val}(\Pi_1) = 1$, and thus $\text{val}(\Pi) = 1$). Hence, $\mathbb{E}_{\langle \Pi_0 \rangle} \left[M_{\Pi_0}^{\mathbf{A}} \right] < 1$, and Definition 3.17 yields that $e_{(\Pi | \neg M_{\Pi}^{\mathbf{A}})}(\lambda, 0) > 0$. By Proposition 3.8, it holds that

$$\begin{aligned}
\text{OPT}_{\mathbf{B}} \left(\Pi | \neg M_{\Pi}^{\mathbf{A}} \right) &\geq \text{OPT}_{\mathbf{B}} \left(\left(\Pi | \neg M_{\Pi}^{\mathbf{A}} \right)_0 \right) \\
&= \text{OPT}_{\mathbf{B}} \left(\Pi_0 | \neg \left(M_{\Pi}^{\mathbf{A}} \right)_0 \right) \\
&= \text{OPT}_{\mathbf{B}} \left(\Pi_0 | \neg M_{\Pi_0}^{\mathbf{A}} \right) \\
&= 1,
\end{aligned}$$

where the second equality follows Proposition 3.19, the third follows the **B**-minimal property of $M_{\Pi}^{\mathbf{A}}$ (Proposition 3.13(2)), and the last equality follows the induction hypothesis. \square

Let $\Pi = (A, B)$ be a protocol in which an optimal adversary playing the role of A biases the outcome towards one with probability one. Lemma 3.20 shows that in the conditional protocol $\Pi_{(B,0)} := \Pi|_{\neg M_{\Pi}^A}$, an optimal adversary playing the role of B can bias the outcome towards zero with probability one. Repeating this procedure with respect to $\Pi_{(B,0)}$ results in the protocol $\Pi_{(A,1)} := \Pi_{(B,0)}|_{\neg M_{\Pi_{(B,0)}}^A}$, in which again an optimal adversary playing the role of A can bias the outcome towards one with probability one. This procedure is stated formally in Definition 3.22.

3.5.2 Sequence of Dominated Measures

Given a protocol (A, B) , we use the simple ordering over the pairs $\{(C, j)\}_{(C,j) \in \{A,B\} \times \mathbb{Z}}$.

Notation 3.21. Let (A, B) be a protocol. For $j \in \mathbb{Z}$ let $\text{pred}(A, j) = (B, j - 1)$ and $\text{pred}(B, j) = (A, j)$, and let succ be the inverse operation of pred (i.e., $\text{succ}(\text{pred}(C, j)) = (C, j)$). For pairs $(C, j), (C', j') \in \{A, B\} \times \mathbb{Z}$, we write

- (C, j) is less than or equal to (C', j') , denoted $(C, j) \preceq (C', j')$, if $\exists \{(C_1, j_1), \dots, (C_n, j_n)\}$ such that $(C, j) = (C_1, j_1)$, $(C', j') = (C_n, j_n)$ and $(C_i, j_i) = \text{pred}(C_{i+1}, j_{i+1})$ for any $i \in [n - 1]$.
- (C, j) is less than (C', j') , denoted $(C, j) \prec (C', j')$, if $(C, j) \preceq (C', j')$ and $(C, j) \neq (C', j')$.

Finally, for $(C, j) \succeq (A, 0)$, let $[(C, j)] := \{(C', j') : (A, 0) \preceq (C', j') \preceq (C, j)\}$.

Definition 3.22. (dominated measures sequence) For a protocol $\Pi = (A, B)$ and $(C, j) \in \{A, B\} \times \mathbb{N}$, the protocol $\Pi_{(C,j)}$ is defined by

$$\Pi_{(C,j)} = \begin{cases} \Pi, & (C, j) = (A, 0); \\ \Pi_{(C',j')=\text{pred}(C,j)}|_{\neg (M_{\Pi_{(C',j')}}^C)}, & \text{otherwise.}^{30} \end{cases}$$

Define the (C, j) dominated measures sequence of Π , denoted (C, j) -DMS(Π), by $\left\{ M_{\Pi_{(C',j')}}^C \right\}_{(C',j') \in [(C,j)]}$. Finally, for $z \in \mathbb{N}$, let $L_{\Pi}^{C,z} \equiv \sum_{j=0}^z M_{\Pi_{(C,j)}}^C \prod_{t=0}^{j-1} (1 - M_{\Pi_{(C,t)}}^C)$.

We show that $L_{\Pi}^{A,z}$ is a measure (i.e., its range is $[0, 1]$) and that its support is a subset of the 1-leaves of Π . We also give an explicit expression for its expected value (analogous to the expected value of M_{Π}^A given in Lemma 3.11).

Lemma 3.23. Let $\Pi = (A, B)$ be a protocol, let $z \in \mathbb{N}$, and let $L_{\Pi}^{A,z}$ be as in Definition 3.22. It holds that

1. $L_{\Pi}^{A,z}$ is a measure over $\mathcal{L}_1(\Pi)$:
 - (a) $L_{\Pi}^{A,z}(\ell) \in [0, 1]$ for every $\ell \in \mathcal{L}(\Pi)$, and
 - (b) $\text{Supp}(L_{\Pi}^{A,z}) \subseteq \mathcal{L}_1(\Pi)$.

³⁰Note that if $E_{\langle \Pi_{(C,j)} \rangle} [M_{\Pi_{(C,j)}}^C] = 1$, Definition 3.17 yields that $\Pi_{\text{succ}(C,j)} = \perp$. In fact, since we defined $\perp|_{\neg M} = \perp$ for any measure M (also in Definition 3.17), it follows that $\Pi_{(C',j')} = \perp$ for any $(C', j') \succ (C, j)$.

2. $\mathbb{E}_{\langle \Pi \rangle} \left[L_{\Pi}^{\mathbf{A},z} \right] = \sum_{j=0}^z \alpha_j \cdot \prod_{t=0}^{j-1} (1 - \beta_t)(1 - \alpha_t)$, where $\alpha_j = 1 - \text{OPT}_{\mathbf{B}}(\Pi_{(\mathbf{A},j)})$, $\beta_j = 1 - \text{OPT}_{\mathbf{A}}(\Pi_{(\mathbf{B},j)})$ and $\text{OPT}_{\mathbf{A}}(\perp) = \text{OPT}_{\mathbf{B}}(\perp) = 1$.

Proof. We prove the above two items separately.

Proof of Item 1. Let $\ell \in \mathcal{L}_0(\Pi)$. Since $M_{\Pi_{(\mathbf{A},j)}}^{\mathbf{A}}(\ell) = 0$ for every $j \in (z)$, it holds that $L_{\Pi}^{\mathbf{A},z}(\ell) = 0$.

Let $\ell \in \mathcal{L}_1(\Pi)$. Since $L_{\Pi}^{\mathbf{A},z}(\ell)$ is a sum of non-negative numbers, it follows that its value is non-negative. It is left to argue that $L_{\Pi}^{\mathbf{A},z}(\ell) \leq 1$. Since $M_{\Pi_{(\mathbf{A},z)}}^{\mathbf{A}}$ is a measure, note that $M_{\Pi_{(\mathbf{A},z)}}^{\mathbf{A}}(\ell) \leq 1$. Thus

$$\begin{aligned}
L_{\Pi}^{\mathbf{A},z}(\ell) &= \sum_{j=0}^z M_{\Pi_{(\mathbf{A},j)}}^{\mathbf{A}}(\ell) \cdot \prod_{t=0}^{j-1} \left(1 - M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}}(\ell) \right) \\
&\leq \prod_{t=0}^{z-1} \left(1 - M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}}(\ell) \right) + \sum_{j=0}^{z-1} M_{\Pi_{(\mathbf{A},j)}}^{\mathbf{A}}(\ell) \cdot \prod_{t=0}^{j-1} \left(1 - M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}}(\ell) \right) \\
&= \left(\sum_{\mathcal{I} \subseteq (z-1)} (-1)^{|\mathcal{I}|} \cdot \prod_{t \in \mathcal{I}} M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}}(\ell) \right) + \sum_{j=0}^{z-1} M_{\Pi_{(\mathbf{A},j)}}^{\mathbf{A}}(\ell) \cdot \left(\sum_{\mathcal{I} \subseteq (j-1)} (-1)^{|\mathcal{I}|} \cdot \prod_{t \in \mathcal{I}} M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}}(\ell) \right) \\
&= \left(\sum_{\mathcal{I} \subseteq (z-1)} (-1)^{|\mathcal{I}|} \cdot \prod_{t \in \mathcal{I}} M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}}(\ell) \right) + \left(\sum_{\emptyset \neq \mathcal{I} \subseteq (z-1)} (-1)^{|\mathcal{I}|+1} \cdot \prod_{t \in \mathcal{I}} M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}}(\ell) \right) \\
&= 1.
\end{aligned}$$

Proof of Item 2. By linearity of expectation, it suffices to prove that

$$\mathbb{E}_{\langle \Pi \rangle} \left[M_{\Pi_{(\mathbf{A},j)}}^{\mathbf{A}} \cdot \prod_{t=0}^{j-1} \left(1 - M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}} \right) \right] = \alpha_j \cdot \prod_{t=0}^{j-1} (1 - \beta_t)(1 - \alpha_t) \quad (15)$$

for any $j \in (z)$. Fix $j \in (z)$. If $\Pi_{(\mathbf{A},j)} = \perp$, then by Definition 3.10 it holds that $M_{\Pi_{(\mathbf{A},j)}}^{\mathbf{A}}$ is the zero measure, and both sides of Equation (15) equal 0.

In the following we assume that $\Pi_{(\mathbf{A},j)} \neq \perp$. We first note that $\mathbb{E}_{\langle \Pi_{(\mathbf{C},t)} \rangle} \left[M_{\Pi_{(\mathbf{C},t)}}^{\mathbf{C}} \right] < 1$ for any $(\mathbf{C}, t) \in [\text{pred}(\mathbf{A}, j)]$ (otherwise, it must be that $\Pi_{(\mathbf{A},j)} = \perp$). Thus, Lemma 3.11 yields that $\alpha_t, \beta_t < 1$ for every $t \in (j-1)$. Hence, recursively applying Proposition 3.18(2) yields that

$$\mathbf{v}_{\langle \Pi_{(\mathbf{A},j)} \rangle}(\ell) = \mathbf{v}_{\Pi}(\ell) \cdot \prod_{t=0}^{j-1} \frac{1 - M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}}(\ell)}{1 - \alpha_t} \cdot \frac{1 - M_{\Pi_{(\mathbf{B},t)}}^{\mathbf{B}}(\ell)}{1 - \beta_t} \quad (16)$$

for every $\ell \in \mathcal{L}(\Pi)$. Moreover, for $\ell \in \text{Supp}(\Pi_{(\mathbf{A},j)})$, i.e., $\mathbf{v}_{\langle \Pi_{(\mathbf{A},j)} \rangle}(\ell) > 0$, we can manipulate Equation (16) to get that

$$\mathbf{v}_{\Pi}(\ell) = \mathbf{v}_{\langle \Pi_{(\mathbf{A},j)} \rangle}(\ell) \cdot \prod_{t=0}^{j-1} \frac{1 - \alpha_t}{1 - M_{\Pi_{(\mathbf{A},t)}}^{\mathbf{A}}(\ell)} \cdot \frac{1 - \beta_t}{1 - M_{\Pi_{(\mathbf{B},t)}}^{\mathbf{B}}(\ell)} \quad (17)$$

for every $\ell \in \text{Supp}(\Pi_{(A,j)})$.

It follows that

$$\begin{aligned}
& \mathbb{E}_{\langle \Pi \rangle} \left[M_{\Pi_{(A,j)}}^A \cdot \prod_{t=0}^{j-1} (1 - M_{\Pi_{(A,t)}}^A) \right] \\
&= \sum_{\ell \in \mathcal{L}(\Pi)} v_{\Pi}(\ell) \cdot \left(M_{\Pi_{(A,j)}}^A(\ell) \cdot \prod_{t=0}^{j-1} (1 - M_{\Pi_{(A,t)}}^A(\ell)) \right) \\
&= \sum_{\ell \in \text{Supp}(\Pi_{(A,j)}) \cap \mathcal{L}_1(\Pi)} v_{\Pi}(\ell) \cdot \left(M_{\Pi_{(A,j)}}^A(\ell) \cdot \prod_{t=0}^{j-1} (1 - M_{\Pi_{(A,t)}}^A(\ell)) \right) \\
&= \sum_{\ell \in \text{Supp}(\Pi_{(A,j)}) \cap \mathcal{L}_1(\Pi)} v_{(\Pi_{(A,j)})}(\ell) \cdot \prod_{t=0}^{j-1} \frac{1 - \alpha_t}{1 - M_{\Pi_{(A,t)}}^A(\ell)} \cdot \frac{1 - \beta_t}{1 - M_{\Pi_{(B,t)}}^B(\ell)} \\
&\quad \cdot \left(M_{\Pi_{(A,j)}}^A(\ell) \cdot \prod_{t=0}^{j-1} (1 - M_{\Pi_{(A,t)}}^A(\ell)) \right) \\
&= \sum_{\ell \in \text{Supp}(\Pi_{(A,j)}) \cap \mathcal{L}_1(\Pi)} v_{(\Pi_{(A,j)})}(\ell) \cdot M_{\Pi_{(A,j)}}^A(\ell) \cdot \prod_{t=0}^{j-1} (1 - \alpha_j)(1 - \beta_j) \\
&= \alpha_j \cdot \prod_{t=0}^{j-1} (1 - \beta_t)(1 - \alpha_t),
\end{aligned}$$

concluding the proof. The second equality follows since Definition 3.10 yields that $M_{\Pi_{(A,j)}}^A(\ell) = 0$ for any $\ell \notin \text{Supp}(\Pi_{(A,j)}) \cap \mathcal{L}_1(\Pi)$, the third equality follows by Equation (17) and the fourth equality follows since $M_{\Pi_{(B,t)}}^B(\ell) = 0$ for every $\ell \in \mathcal{L}_1(\Pi)$ and $t \in (j-1)$. \square

Example 3.24. *Once again we consider the protocol Π from Figure 2a. In Figure 3 we present the conditional protocol $\Pi_{(B,0)} = \Pi | \neg M_{\Pi}^A$, namely the protocol derived when protocol Π is conditioned not to “hit” the A-dominated measure of Π . We would like to highlight some points regarding this conditional protocol.*

The first point we note is the changes in the edge distribution. Now consider the root of Π_0 (i.e., the node 0). According to the calculations in Figure 2b, it holds that $\mathbb{E}_{\langle \Pi_0 \rangle} [M_{\Pi}^A] = M_{\Pi}^A(00) = 1$ and that $\mathbb{E}_{\langle \Pi_0 \rangle} [M_{\Pi}^A] = \alpha_0$. Hence, Definition 3.17 yields that

$$\begin{aligned}
e_{(\Pi | \neg M_{\Pi}^A)}(0, 00) &= \alpha_0 \cdot \frac{1 - \mathbb{E}_{\langle \Pi_0 \rangle} [M_{\Pi}^A]}{1 - \mathbb{E}_{\langle \Pi_0 \rangle} [M_{\Pi}^A]} \\
&= \alpha_0 \cdot \frac{0}{1 - \alpha_0} \\
&= 0.
\end{aligned}$$

Note that the above change makes the leaf 00 inaccessible in $\Pi_{(B,0)}$. This occurs since $M_{\Pi}^A(00) = 1$ and follows Proposition 3.18. Similar calculations yield the changes in the distribution of the edges leaving the root of Π_1 (i.e., the node 1).




Figure 3: The conditional protocol $\Pi_{(B,0)} = \Pi|_{\neg M_{\Pi}^A}$ of Π from Figure 2a. Dashed edges are such that their distribution has changed. Note that due to this change, the leaf 00 (the leftmost leaf, marked by a thick border) is *inaccessible* in $\Pi_{(B,0)}$. The B-dominated measure of $\Pi_{(B,0)}$ assigns a value of 1 to the leaf 010, and value of 0 to all other leaves.

The second point we note is that the conditional protocol is in fact a protocol. Namely, for every node, the sum of the probabilities of the edges leaving it is one. This is easily seen from Figure 3 and again follows from Proposition 3.18.

The third point we note is that the edge distribution of the root of Π does not change at all. This follows from Definition 3.17 and the fact that

$$\mathbb{E}_{\langle \Pi_0 \rangle} [M_{\Pi}^A] = \mathbb{E}_{\langle \Pi_1 \rangle} [M_{\Pi}^A] = \mathbb{E}_{\langle \Pi \rangle} [M_{\Pi}^A] = \alpha_0.$$

The fourth point we note is that in the conditional protocol, an optimal valid attacker playing the role of B can bias the outcome towards zero with probability one. Namely, $\text{OPT}_B(\Pi|_{\neg M_{\Pi}^A}) = 1$. Such an attacker will send 0 as the first message, after which A must send 1 as the next message, and then the attacker will send 0. The outcome of this interaction is the value of the leaf 010, which is 0. This follows from Lemma 3.20.

Using dominated measure sequences, we manage to give an improved bound for the success probability of the recursive biased-continuation attacks (comparing to the bound of Lemma 3.16, which uses a single dominated measure). The improved analysis yields that a constant number of recursion calls of the biased-continuation attack is successful in biasing the protocol to an arbitrary constant close to either 0 or 1.

3.6 Improved Analysis Using Alternating Dominated Measures

We are finally ready to state two main lemmas, whose proofs – given in the next two sections – are the main technical contribution of Section 3, and then show how to use them to prove Theorem 3.3.

The first lemma is analogous to Lemma 3.14, but applied on the sequence of the dominated measures, and not just on a single dominated measure.

Lemma 3.25. *For a protocol $\Pi = (A, B)$ with $\text{val}(\Pi) > 0$ and $z \in \mathbb{N}$, it holds that*

$$\text{val}(A^{(k)}, B) \geq \mathbb{E}_{\langle A^{(k)}, B \rangle} [L_{\Pi}^{A,z}] \geq \frac{\mathbb{E}_{\langle \Pi \rangle} [L_{\Pi}^{A,z}]}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)} \cdot \left(1 - \sum_{j=0}^{z-1} \beta_j \right)^k$$

for every $k \in \mathbb{N}$, where $\beta_j = 1 - \text{OPT}_A(\Pi_{(B,j)})$, letting $\text{OPT}_A(\perp) = 1$.

The above states that the recursive biased-continuation attacker biases the outcome of the protocol by a bound similar to that given in Lemma 3.14, but applied with respect to $L_{\Pi}^{A,z}$, instead of M_{Π}^A in Lemma 3.14. This is helpful since the expected value of $L_{\Pi}^{A,z}$ is strictly larger than that of M_{Π}^A . However, since $L_{\Pi}^{A,z}$ is defined with respect to a sequence of conditional protocols, we must “pay” the term $\left(1 - \sum_{j=0}^{z-1} \beta_j\right)^k$ in order to get this bound in the original protocol.

The following lemma states that Lemma 3.25 provides a sufficient bound. Specifically, it shows that if we take a long enough sequence of conditional protocols, the expected value of the measure $L_{\Pi}^{A,z}$ will be sufficiently large, while the payment term mentioned above will be kept sufficiently small.

Lemma 3.26. *Let $\Pi = (A, B)$ be a protocol. Then for every $c \in (0, \frac{1}{2}]$ there exists $z = z(c, \Pi) \in \mathbb{N}$ (possibly exponential large) such that:*

1. $\mathbb{E}_{(\Pi)} \left[L_{\Pi}^{A,z} \right] \geq c \cdot (1 - 2c)$ and $\sum_{j=0}^{z-1} \beta_j < c$; or
2. $\mathbb{E}_{(\Pi)} \left[L_{\Pi}^{B,z} \right] \geq c \cdot (1 - 2c)$ and $\sum_{j=0}^z \alpha_j < c$,

where $\alpha_j = 1 - \text{OPT}_B(\Pi_{(A,j)})$ and $\beta_j = 1 - \text{OPT}_A(\Pi_{(B,j)})$.

To derive Theorem 3.3, we take a sequence of the dominated measures that is long enough so that its accumulated weight will be sufficiently large. Furthermore, the weight of the dominated measures that precede the final dominated measure in the sequence is small (otherwise, we would have taken a shorter sequence), so the parties are “missing” these measures with high probability. The formal proof of Theorem 3.3 is given next, and the proofs of Lemmas 3.25 and 3.26 are given in Sections 3.7 and 3.8 respectively.

3.6.1 Proving Theorem 3.3

Proof of Theorem 3.3. If $\text{val}(\Pi) = 0$, Theorem 3.3 trivially holds. Assume that $\text{val}(\Pi) > 0$, let z be the minimum integer guaranteed by Lemma 3.26 for $c = \varepsilon/2$, and let $\kappa = \left\lceil \frac{\log(\frac{2}{\varepsilon})}{\log(\frac{1-\varepsilon/2}{1-\varepsilon})} \right\rceil$.

If z satisfies Item 1 of Lemma 3.26, assume towards a contradiction that $\text{val}(A^{(\kappa)}, B) \leq 1 - \varepsilon$. Lemma 3.25 yields that

$$\begin{aligned} \text{val}(A^{(\kappa)}, B) &\geq \frac{\mathbb{E}_{(\Pi)} \left[L_{\Pi}^{A,z} \right]}{\prod_{i=0}^{\kappa-1} \text{val}(A^{(i)}, B)} \cdot \left(1 - \sum_{j=0}^{z-1} \beta_j \right)^{\kappa} \\ &> \frac{\varepsilon(1-\varepsilon)}{2} \cdot \left(\frac{1-\varepsilon/2}{1-\varepsilon} \right)^{\kappa} \\ &\geq 1 - \varepsilon, \end{aligned}$$

and a contradiction is derived.

If z satisfies Item 2 of Lemma 3.26, an analogous argument to the above yields that $\text{val}(A, B^{(\kappa)}) \leq \varepsilon$. □

3.7 Proving Lemma 3.25

The proof of Lemma 3.25 is an easy implication of Lemma 3.23 and the following key lemma, defined with respect to sequences of *submeasures* of the dominated measure.

Definition 3.27. (*dominated submeasure sequence*) For a protocol $\Pi = (A, B)$, a pair $(C^*, j^*) \in \{A, B\} \times \mathbb{N}$ and $\boldsymbol{\eta} = \{\eta_{(C,j)} \in [0, 1]\}_{(C,j) \in [(C^*, j^*)]}$, define the protocol $\widehat{\Pi}_{(C,j)}^\boldsymbol{\eta}$ by

$$\widehat{\Pi}_{(C,j)}^\boldsymbol{\eta} := \begin{cases} \Pi, & (C, j) = (A, 0); \\ \widehat{\Pi}_{(C',j')^\text{pred}(C,j)}^\boldsymbol{\eta} \upharpoonright \neg \left(\widehat{M}_{(C',j')}^{\Pi, \boldsymbol{\eta}} \right), & \text{otherwise.} \end{cases},$$

where $\widehat{M}_{(C',j')}^{\Pi, \boldsymbol{\eta}} \equiv \eta_{(C',j')} \cdot M_{\Pi_{(C',j')}}^{C'}$. For $(C, j) \in [(C^*, j^*)]$, define the $(C, j, \boldsymbol{\eta})$ -dominated measure sequence of Π , denoted $(C, j, \boldsymbol{\eta})$ -DMS (Π) , as $\left\{ \widehat{M}_{(C',j')}^{\Pi, \boldsymbol{\eta}} \right\}_{(C',j') \in [(C,j)]}$, and let $\widehat{\mu}_{(C,j)}^{\Pi, \boldsymbol{\eta}} = \mathbb{E}_{\langle \widehat{\Pi}_{(C,j)}^\boldsymbol{\eta} \rangle} \left[\widehat{M}_{(C,j)}^{\Pi, \boldsymbol{\eta}} \right]$.³¹

Finally, let $\widehat{L}_\Pi^{C, \boldsymbol{\eta}} \equiv \sum_{j: (C,j) \in [(C^*, j^*)]} \widehat{M}_{(C,j)}^{\Pi, \boldsymbol{\eta}} \cdot \prod_{t=0}^{j-1} (1 - \widehat{M}_{(C,t)}^{\Pi, \boldsymbol{\eta}})$.

Lemma 3.28. Let $\Pi = (A, B)$ be a protocol with $\text{val}(\Pi) > 0$, let $z \in \mathbb{N}$ and let $\boldsymbol{\eta} = \{\eta_{(C,j)} \in [0, 1]\}_{(C,j) \in [(A, z)]}$. For $j \in (z)$, let $\alpha_j = \widehat{\mu}_{(A,j)}^{\Pi, \boldsymbol{\eta}}$, and for $j \in (z-1)$, let $\beta_j = \widehat{\mu}_{(B,j)}^{\Pi, \boldsymbol{\eta}}$. Then

$$\mathbb{E}_{\langle A^{(k)}, B \rangle} \left[\widehat{L}_\Pi^{A, \boldsymbol{\eta}} \right] \geq \frac{\sum_{j=0}^z \alpha_j \cdot \prod_{t=0}^{j-1} (1 - \beta_t)^{k+1} (1 - \alpha_t)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)}$$

for any positive $k \in \mathbb{N}$.

The proof of Lemma 3.28 is given below, but we first use it to prove Lemma 3.25.

Proof of Lemma 3.25. Let $\eta_{(C,j)} = 1$ for every $(C, j) \in [(A, z)]$ and let $\boldsymbol{\eta} = \{\eta_{(C,j)}\}_{(C,j) \in [(A, z)]}$. It follows that $\widehat{L}_\Pi^{A, \boldsymbol{\eta}} \equiv L_\Pi^{A, z}$. Applying Lemma 3.28 yields that

$$\mathbb{E}_{\langle A^{(k)}, B \rangle} \left[L_\Pi^{A, z} \right] \geq \frac{\sum_{j=0}^z \alpha_j \cdot \prod_{t=0}^{j-1} (1 - \beta_t)^{k+1} (1 - \alpha_t)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)}, \quad (18)$$

where $\alpha_j = \widehat{\mu}_{(A,j)}^{\Pi, \boldsymbol{\eta}}$ and $\beta_j = \widehat{\mu}_{(B,j)}^{\Pi, \boldsymbol{\eta}}$. Multiplying the j 'th summand of the right-hand side of Equation (18) by $\prod_{t=j}^{z-1} (1 - \beta_j)^k \leq 1$ yields that

$$\begin{aligned} \mathbb{E}_{\langle A^{(k)}, B \rangle} \left[L_\Pi^{A, z} \right] &\geq \frac{\sum_{j=0}^z \alpha_j \cdot \prod_{t=0}^{j-1} (1 - \beta_t) (1 - \alpha_t)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)} \cdot \prod_{t=0}^{z-1} (1 - \beta_t)^k \\ &\geq \frac{\sum_{j=0}^z \alpha_j \cdot \prod_{t=0}^{j-1} (1 - \beta_t) (1 - \alpha_t)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)} \cdot \left(1 - \sum_{t=0}^{z-1} \beta_t \right)^k, \end{aligned} \quad (19)$$

³¹Note that for $\boldsymbol{\eta} = (1, 1, 1, \dots, 1)$, Definition 3.27 coincides with Definition 3.22.

where the second inequality follows since $\beta_j \geq 0$ and $(1-x)(1-y) \geq 1-(x+y)$ for any $x, y \geq 0$. By Lemma 3.11 and the definition of $\boldsymbol{\eta}$ it follows that $\widehat{\mu}_{(\mathbf{A},j)}^{\Pi,\boldsymbol{\eta}} = 1 - \text{OPT}_{\mathbf{B}}(\Pi_{(\mathbf{A},j)})$ and $\widehat{\mu}_{(\mathbf{B},j)}^{\Pi,\boldsymbol{\eta}} = 1 - \text{OPT}_{\mathbf{A}}(\Pi_{(\mathbf{B},j)})$. Hence, plugging Lemma 3.23 into Equation (19) yields that

$$\mathbb{E}_{\langle \mathbf{A}^{(k)}, \mathbf{B} \rangle} \left[L_{\Pi}^{\mathbf{A},z} \right] \geq \frac{\mathbb{E}_{\langle \Pi \rangle} \left[L_{\Pi}^{\mathbf{A},z} \right]}{\prod_{i=0}^{k-1} \text{val}(\mathbf{A}^{(i)}, \mathbf{B})} \cdot \left(1 - \sum_{t=0}^{z-1} \beta_t \right)^k. \quad (20)$$

Finally, the proof is concluded, since by Lemma 3.23 and Fact 2.5 it immediately follows that $\text{val}(\mathbf{A}^{(k)}, \mathbf{B}) \geq \mathbb{E}_{\langle \mathbf{A}^{(k)}, \mathbf{B} \rangle} \left[L_{\Pi}^{\mathbf{A},z} \right]$. \square

3.7.1 Proving Lemma 3.28

Proof of Lemma 3.28. In the following we fix a protocol Π , real vector $\boldsymbol{\eta} = \{\eta_{(\mathbf{C},j)}\}_{(\mathbf{C},j) \in [(\mathbf{A},z)]}$ and a positive integer k . We also assume for simplicity that $\widehat{\Pi}_{(\mathbf{A},z)}^{\boldsymbol{\eta}}$ is not the undefined protocol, i.e., $\widehat{\Pi}_{(\mathbf{A},z)}^{\boldsymbol{\eta}} \neq \perp$.³² The proof is by induction on the round complexity of Π .

Base case. Assume $\text{round}(\Pi) = 0$ and let ℓ be the only node in $\mathcal{T}(\Pi)$. For $j \in (z)$, Definition 3.27 yields that $\chi_{\widehat{\Pi}_{(\mathbf{A},j)}^{\boldsymbol{\eta}}}(\ell) = \chi_{\Pi}(\ell) = 1$, where the last equality holds since, by assumption, $\text{val}(\Pi) > 0$. It follows Definition 3.10 that $M_{\widehat{\Pi}_{(\mathbf{A},j)}^{\boldsymbol{\eta}}}^{\mathbf{A}}(\ell) = 1$ and Definition 3.27 that $\widehat{M}_{(\mathbf{A},j)}^{\Pi,\boldsymbol{\eta}}(\ell) = \eta_{(\mathbf{A},j)}$. Hence, it holds that $\alpha_j = \eta_{(\mathbf{A},j)}$. Similarly, for $j \in (z-1)$ it holds that $\widehat{M}_{(\mathbf{B},j)}^{\Pi,\boldsymbol{\eta}}(\ell) = 0$ and thus $\beta_j = 0$. Clearly, $(\mathbf{A}^{(k)}, \mathbf{B}) = \Pi$ and $\text{val}(\mathbf{A}^{(i)}, \mathbf{B}) = 1$ for every $i \in [k-1]$. We conclude that

$$\begin{aligned} \mathbb{E}_{\langle \mathbf{A}^{(k)}, \mathbf{B} \rangle} \left[\widehat{L}_{\mathbf{A}}^{\Pi,\boldsymbol{\eta}} \right] &= \mathbb{E}_{\langle \Pi \rangle} \left[\widehat{L}_{\mathbf{A}}^{\Pi,\boldsymbol{\eta}} \right] \\ &= \sum_{j=0}^z \widehat{M}_{(\mathbf{A},j)}^{\Pi,\boldsymbol{\eta}}(\ell) \cdot \prod_{t=0}^{j-1} \left(1 - \widehat{M}_{(\mathbf{A},t)}^{\Pi,\boldsymbol{\eta}}(\ell) \right) \\ &= \sum_{j=0}^z \eta_{(\mathbf{A},j)} \cdot \prod_{t=0}^{j-1} (1 - \eta_{(\mathbf{A},t)}) \\ &= \sum_{j=0}^z \alpha_j \cdot \prod_{t=0}^{j-1} (1 - \alpha_t) \\ &= \frac{\sum_{j=0}^z \alpha_j \prod_{t=0}^{j-1} (1 - \beta_t)^{k+1} (1 - \alpha_t)}{\prod_{i=0}^{k-1} \text{val}(\mathbf{A}^{(i)}, \mathbf{B})}. \end{aligned}$$

Induction step. Assume the lemma holds for m -round protocols and that $\text{round}(\Pi) = m+1$. We prove it by the following steps: (1) we define two real vectors $\boldsymbol{\eta}_0$ and $\boldsymbol{\eta}_1$ such that the restriction

³²If this assumption does not hold, let $z' \in (z-1)$ be the largest index such that $\widehat{\Pi}_{(\mathbf{A},z')}^{\boldsymbol{\eta}} \neq \perp$, and let $\boldsymbol{\eta}' = \{\eta_{(\mathbf{C},j)}\}_{(\mathbf{C},j) \in [(\mathbf{A},z')]}$. It follows from Definition 3.10 that $\widehat{M}_{(\mathbf{A},j)}^{\Pi,\boldsymbol{\eta}'}$ is the zero measure for any $z' < j \leq z$, and thus $\widehat{L}_{\mathbf{A}}^{\Pi,\boldsymbol{\eta}'} \equiv \widehat{L}_{\mathbf{A}}^{\Pi,\boldsymbol{\eta}}$. Moreover, the fact that $\alpha_j = 0$ for any $z' < j \leq z$ suffices to validate the assumption.

of $\widehat{L}_A^{\Pi, \eta}$ to Π_0 and Π_1 is equal to $\widehat{L}_A^{\Pi_0, \eta_0}$ and $\widehat{L}_A^{\Pi_1, \eta_1}$ respectively; (2) we apply the induction hypothesis on the two latter measures; (3) if \mathbf{A} controls $\text{root}(\Pi)$, we use the properties of $\mathbf{A}^{(k)}$ – as stated in Claim 3.4 – to derive the lemma, whereas if \mathbf{B} controls $\text{root}(\Pi)$, we derive the lemma from Lemma 2.17.

All claims given in the context of this proof are proven in Section 3.7.2. We defer handling the case that $e_\Pi(\lambda, b) \in \{0, 1\}$ for some $b \in \{0, 1\}$ (see the end of this proof) and assume for now that $e_\Pi(\lambda, 0), e_\Pi(\lambda, 1) \in (0, 1)$. The real vectors η_0 and η_1 are defined as follows.

Definition 3.29. Let $\eta_{\mathbf{b}} = \left\{ \eta_{(\mathbf{C}, j)}^b \right\}_{(\mathbf{C}, j) \in [(A, z)]}$, where for $(\mathbf{C}, j) \in [(A, z)]$ and $b \in \{0, 1\}$, let

$$\eta_{(\mathbf{C}, j)}^b = \begin{cases} 0 & e_{\widehat{\Pi}_{(\mathbf{C}, j)}^\eta}(\lambda, b) = 0; \\ \eta_{(\mathbf{C}, j)} & e_{\widehat{\Pi}_{(\mathbf{C}, j)}^\eta}(\lambda, b) = 1; \\ \eta_{(\mathbf{C}, j)} & e_{\widehat{\Pi}_{(\mathbf{C}, j)}^\eta}(\lambda, b) \notin \{0, 1\} \wedge (\mathbf{C} \text{ controls } \text{root}(\Pi) \vee \text{Smaller}_{\widehat{\Pi}_{(\mathbf{C}, j)}^\eta}(b)); \\ \frac{\xi_{(\mathbf{C}, j)}^{1-b}}{\xi_{(\mathbf{C}, j)}^b} \cdot \eta_{(\mathbf{C}, j)} & \text{otherwise;} \end{cases} ,$$

where $\xi_{(\mathbf{C}, j)}^b = \mathbb{E} \langle (\widehat{\Pi}_{(\mathbf{C}, j)}^\eta)_b \rangle \left[M_{(\widehat{\Pi}_{(\mathbf{C}, j)}^\eta)_b}^{\mathbf{C}} \right]$ and $\text{Smaller}_{\widehat{\Pi}_{(\mathbf{C}, j)}^\eta}(b) = 1$ if $\xi_{(\mathbf{C}, j)}^b \leq \xi_{(\mathbf{C}, j)}^{1-b}$.³³

Given the real vector $\eta_{\mathbf{b}}$, consider the dominated submeasure sequence $\eta_{\mathbf{b}}$ induces on the subprotocol Π_b . At first glance, the relation of this submeasure sequence to the dominated submeasure sequence η induces on Π , is unclear; nonetheless, we manage to prove the following key observation.

Claim 3.30. It holds that $\widehat{L}_A^{\Pi_b, \eta_{\mathbf{b}}} \equiv \left(\widehat{L}_A^{\Pi, \eta} \right)_b$ for both $b \in \{0, 1\}$.

Namely, taking $(A, z, \eta_{\mathbf{b}})$ -DMS (Π_b) – the dominated submeasures defined with respect to Π_b and $\eta_{\mathbf{b}}$ – and constructing from it the measure $\widehat{L}_A^{\Pi_b, \eta_{\mathbf{b}}}$, results in the same measure as taking (A, z, η) -DMS (Π) – the dominated submeasures defined with respect to Π and η – and constructing from it the measure $\widehat{L}_A^{\Pi, \eta}$ while restricting the latter to Π_b .

Given the above fact, we can use our induction hypothesis on the subprotocols Π_0 and Π_1 with respect to the real vectors η_0 and η_1 , respectively. For $b \in \{0, 1\}$ and $j \in (z)$, let $\alpha_j^b := \mu_{(A, j)}^{\Pi_b, \eta_{\mathbf{b}}}$ ($:= \mathbb{E} \langle (\widehat{\Pi}_b)_{(A, j)}^{\eta_{\mathbf{b}}} \rangle \left[\widehat{M}_{(A, j)}^{\Pi_b, \eta_{\mathbf{b}}} \right]$), and for $j \in (z-1)$ let $\beta_j^b := \mu_{(B, j)}^{\Pi_b, \eta_{\mathbf{b}}}$. Assuming that $\text{val}(\Pi_1) > 0$, then

$$\mathbb{E} \langle (\mathbf{A}^{(k), \mathbf{B}})_1 \rangle \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_1 \right] = \mathbb{E} \langle \mathbf{A}_{\Pi_1}^{(k), \mathbf{B}_{\Pi_1}} \rangle \left[\widehat{L}_A^{\Pi_1, \eta_1} \right] \geq \frac{\sum_{j=0}^z \alpha_j^1 \prod_{t=0}^{j-1} (1 - \beta_t^1)^{k+1} (1 - \alpha_t^1)}{\prod_{i=0}^{k-1} \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_1)}. \quad (21)$$

where the equality holds by Proposition 3.5 and Claim 3.30, and the inequality by the induction hypothesis. Similarly, if $\text{val}(\Pi_0) > 1$, then

$$\mathbb{E} \langle (\mathbf{A}^{(k), \mathbf{B}})_0 \rangle \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_0 \right] = \mathbb{E} \langle \mathbf{A}_{\Pi_0}^{(k), \mathbf{B}_{\Pi_0}} \rangle \left[\widehat{L}_A^{\Pi_0, \eta_0} \right] \geq \frac{\sum_{j=0}^z \alpha_j^0 \prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1} (1 - \alpha_t^0)}{\prod_{i=0}^{k-1} \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_0)}. \quad (22)$$

³³Note that the definition of η^b follows the same lines of the definition of the dominated measure (given in Definition 3.10).

In the following we use the fact that the dominated submeasure sequence of one of the subprotocols is at least as long as the submeasure sequence of the protocol itself. Specifically, we show the following.

Definition 3.31. For $b \in \{0, 1\}$, let $z^b = \min \left\{ \left\{ j \in (z) : \alpha_j^b = 1 \vee \beta_j^b = 1 \right\} \cup \{z\} \right\}$.

Assuming without loss of generality (and throughout the proof of the lemma) that $z^1 \leq z^0$, we have the following claim (proven in Section 3.7.2).

Claim 3.32. Assume that $z^1 \leq z^0$, then $z^0 = z$.

We are now ready to prove the lemma by separately considering which party controls the root of Π .

A controls $\text{root}(\Pi)$ and $\text{val}(\Pi_0), \text{val}(\Pi_1) > 0$. Under these assumptions, we can apply the induction hypothesis on both subtrees (namely, we can use Equations (21) and (22)). Let $p = e_\Pi(\lambda, 0)$. Compute

$$\begin{aligned}
& \mathbb{E}_{\langle A^{(k)}, B \rangle} \left[\widehat{L}_A^{\Pi, \eta} \right] \tag{23} \\
&= e_{\langle A^{(k)}, B \rangle}(\lambda, 0) \cdot \mathbb{E}_{\langle (A^{(k)}, B)_0 \rangle} \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_0 \right] + e_{\langle A^{(k)}, B \rangle}(\lambda, 1) \cdot \mathbb{E}_{\langle (A^{(k)}, B)_1 \rangle} \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_1 \right] \\
&= p \cdot \frac{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_0)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)} \cdot \mathbb{E}_{\langle (A^{(k)}, B)_0 \rangle} \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_0 \right] \\
&\quad + (1-p) \cdot \frac{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_1)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)} \cdot \mathbb{E}_{\langle (A^{(k)}, B)_1 \rangle} \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_1 \right] \\
&\geq p \cdot \frac{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_0)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)} \cdot \frac{\sum_{j=0}^z \alpha_j^0 \prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1} (1 - \alpha_t^0)}{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_0)} \\
&\quad + (1-p) \cdot \frac{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_1)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)} \cdot \frac{\sum_{j=0}^z \alpha_j^1 \prod_{t=0}^{j-1} (1 - \beta_t^1)^{i+1} (1 - \alpha_t^1)}{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_1)} \\
&= \frac{p \cdot \left(\sum_{j=0}^z \alpha_j^0 \prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1} (1 - \alpha_t^0) \right)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)} + \frac{(1-p) \cdot \left(\sum_{j=0}^z \alpha_j^1 \prod_{t=0}^{j-1} (1 - \beta_t^1)^{k+1} (1 - \alpha_t^1) \right)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)},
\end{aligned}$$

where the second equality follows Claim 3.4 and the third inequality follows Equations (21) and (22).

Our next step is to establish a connection between the above $\left\{ \alpha_j^0, \alpha_j^1 \right\}_{j \in (z)}$ and $\left\{ \beta_j^0, \beta_j^1 \right\}_{j \in (z-1)}$ to $\left\{ \alpha_j \right\}_{j \in (z)}$ and $\left\{ \beta_j \right\}_{j \in (z-1)}$ (appearing in the lemma's statement). We prove the following claims.

Claim 3.33. If A controls $\text{root}(\Pi)$, it holds that $\beta_j^0 = \beta_j$ for every $j \in (z-1)$ and $\beta_j^1 = \beta_j$ for every $j \in (z^1-1)$.

It is a direct implication of Proposition 3.13 that $\beta_j^0 = \beta_j^1 = \beta_j$ for $j \in (z^1-1)$. Moreover, $\beta_j^0 = \beta_j$ for every $z^1 \leq j \leq z-1$. The latter is harder to grasp without the technical proof of the claim, which is provided in Section 3.7.2.

Claim 3.34. *If A controls $\text{root}(\Pi)$ and $z^1 < z$, it holds that $\alpha_{z^1}^1 = 1$.*

By Claim 3.33 it follows that as long as an undefined protocol was not reached in one of the subprotocols, then $\beta_j^0 = \beta_j^1 = \beta_j$. Assuming that $z^1 < z$ and $\beta_{z^1}^1 = 1$, it would have followed that $\beta_{z^1} = 1$, and an undefined protocol is reached in the original protocol before z , a contradiction to our assumption. (Again, see Section 3.7.2 for the formal proof.)

Claims 3.33 and 3.34 and Equation (23) yield that

$$\mathbb{E}_{\langle A^{(k)}, B \rangle} \left[\widehat{L}_A^{\Pi, \eta} \right] \geq \frac{\sum_{j=0}^z \prod_{t=0}^{j-1} (1 - \beta_t)^{k+1} \left(p \cdot \alpha_j^0 \prod_{t=0}^{j-1} (1 - \alpha_t^0) + (1 - p) \cdot \alpha_j^1 \cdot \prod_{t=0}^{j-1} (1 - \alpha_t^1) \right)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)}. \quad (24)$$

The proof of this case is concluded by plugging the next claim into Equation (24).

Claim 3.35. *If A controls $\text{root}(\Pi)$ it holds that*

$$\alpha_j \cdot \prod_{t=0}^{j-1} (1 - \alpha_t) = p \cdot \alpha_j^0 \cdot \prod_{t=0}^{j-1} (1 - \alpha_t^0) + (1 - p) \cdot \alpha_j^1 \cdot \prod_{t=1}^{j-1} (1 - \alpha_t^1)$$

for any $j \in (z)$.

Claim 3.35 is proven in Section 3.7.2, but informally it holds since the probability of visiting the left-hand [resp., right-hand] subprotocol in the conditional protocol $\widehat{\Pi}_{(A,j)}^\eta$ (in which α_j is defined) is $p \cdot \prod_{t=0}^{j-1} (1 - \alpha_t^0) / \prod_{t=0}^{j-1} (1 - \alpha_t)$ [resp., $(1 - p) \cdot \prod_{t=0}^{j-1} (1 - \alpha_t^1) / \prod_{t=0}^{j-1} (1 - \alpha_t)$]. Since α_j is defined to be the expected value of some measure in the above conditional protocol, its value is a linear combination of α_j^0 and α_j^1 , with the coefficient being the above probabilities.

A controls $\text{root}(\Pi)$ and $\text{val}(\Pi_0) > \text{val}(\Pi_1) = 0$. Under these assumptions, we can still use the induction hypothesis for the left-hand subprotocol Π_0 , where for right-hand subprotocol Π_1 , we argue the following.

Claim 3.36. *If $\text{val}(\Pi_1) = 0$, it holds that $\left(\widehat{L}_A^{\Pi, \eta} \right)_1 \equiv 0$.³⁴*

Claim 3.36 holds since according to Claim 3.30 we can simply argue that $\widehat{L}_A^{\Pi_1, \eta_1}$ is the zero measure, and this holds since the latter measure is a combination of A-dominated measures, all of which are the zero measure in a zero-value protocol.

Using Claim 3.36, similar computations to the ones in Equation (23) yield that

$$\begin{aligned} & \mathbb{E}_{\langle A^{(k)}, B \rangle} \left[\widehat{L}_A^{\Pi, \eta} \right] \quad (25) \\ &= e_{\langle A^{(k)}, B \rangle}(\lambda, 0) \cdot \mathbb{E}_{\langle (A^{(k)}, B)_0 \rangle} \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_0 \right] + e_{\langle A^{(k)}, B \rangle}(\lambda, 1) \cdot \mathbb{E}_{\langle (A^{(k)}, B)_1 \rangle} \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_1 \right] \\ &\geq p \cdot \frac{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_0)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)} \cdot \frac{\sum_{j=0}^z \alpha_j^0 \prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1} (1 - \alpha_t^0)}{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_0)} \\ &= \frac{p \cdot \left(\sum_{j=0}^z \alpha_j^0 \prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1} (1 - \alpha_t^0) \right)}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)}. \end{aligned}$$

³⁴That is, $\left(\widehat{L}_A^{\Pi, \eta} \right)_1$ is the zero measure.

Using a similar argument to that of Equation (24), combining Claim 3.33 and Equation (25) yields that

$$\mathbb{E}_{\langle A^{(k)}, B \rangle} \left[\widehat{L}_A^{\Pi, \eta} \right] \geq \frac{\sum_{j=0}^z \prod_{t=0}^{j-1} (1 - \beta_t)^{k+1} \left[p \cdot \alpha_j^0 \prod_{t=0}^{j-1} (1 - \alpha_t^0) \right]}{\prod_{i=0}^{k-1} \text{val}(A^{(i)}, B)}. \quad (26)$$

The proof of this case is concluded by plugging the next claim (proven in Section 3.7.2) into Claim 3.35, and plugging the result into Equation (26).

Claim 3.37. *If $\text{val}(\Pi_1) = 0$, it holds that $\alpha_j^1 = 0$ for every $j \in (z)$.*

A controls $\text{root}(\Pi)$ and $\text{val}(\Pi_1) > \text{val}(\Pi_0) = 0$. The proof of the lemma under these assumptions is analogous to the previous case.

We have concluded the proof for cases in which A controls $\text{root}(\Pi)$, and now proceed to prove the cases in which B controls $\text{root}(\Pi)$. Roughly speaking, A and B switched roles, and claims true before regarding β_j are now true for α_j , and vice versa. Moreover, the analysis above relies on the probabilities that the recursive biased-continuation attacker visits the subprotocols Π_0 and Π_1 when it plays the role of A and controls $\text{root}(\Pi)$. When B controls $\text{root}(\Pi)$, however, these probabilities do not change (namely, they remain p and $1 - p$ respectively). To overcome this difficulty we use a convex type argument stated in Lemma 2.17.

B controls $\text{root}(\Pi)$ and $\text{val}(\Pi_0), \text{val}(\Pi_1) > 0$. In this case Equations (21) and (22) hold.

Compute

$$\begin{aligned} & \mathbb{E}_{\langle A^{(k)}, B \rangle} \left[\widehat{L}_A^{\Pi, \eta} \right] \quad (27) \\ &= p \cdot \mathbb{E}_{\langle (A^{(k)}, B)_0 \rangle} \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_0 \right] + (1 - p) \cdot \mathbb{E}_{\langle (A^{(k)}, B)_1 \rangle} \left[\left(\widehat{L}_A^{\Pi, \eta} \right)_1 \right] \\ &\geq p \cdot \frac{\sum_{j=0}^z \alpha_j^0 \prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1} (1 - \alpha_t^0)}{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_0)} + (1 - p) \cdot \frac{\sum_{j=0}^z \alpha_j^1 \prod_{t=0}^{j-1} (1 - \beta_t^1)^{k+1} (1 - \alpha_t^1)}{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_1)}, \end{aligned}$$

where the inequality follows from Equations (21) and (22). If B controls $\text{root}(\Pi)$, we can prove the next claims (proven in Section 3.7.2), analogous to Claims 3.33 and 3.34.

Claim 3.38. *If B controls $\text{root}(\Pi)$, it holds that $\alpha_j^0 = \alpha_j$ for every $j \in (z)$ and that $\alpha_j^1 = \alpha_j$ for every $j \in (z^1)$.*

Claim 3.39. *If B controls $\text{root}(\Pi)$ and $z^1 < z$, it holds that $\beta_{z^1}^1 = 1$.*

Claim 3.38 and Equation (27) yield that

$$\begin{aligned} & \mathbb{E}_{\langle A^{(k)}, B \rangle} \left[\widehat{L}_A^{\Pi, \eta} \right] \quad (28) \\ &\geq \sum_{j=0}^z \alpha_j \prod_{t=0}^{j-1} (1 - \alpha_t) \left(p \cdot \frac{\prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1}}{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_0)} + (1 - p) \cdot \frac{\prod_{t=0}^{j-1} (1 - \beta_t^1)^{k+1}}{\prod_{i=0}^{k-1} \text{val}((A^{(i)}, B)_1)} \right). \end{aligned}$$

Applying the convex type inequality given in Lemma 2.17 for each summand in the right-hand side of Equation (28) with respect to $x = \prod_{t=0}^{j-1} (1 - \beta_t^0)$, $y = \prod_{t=0}^{j-1} (1 - \beta_t^1)$, $a_i = \text{val}(\mathbf{A}^{(i-1)}, \mathbf{B}_0)$, $b_i = \text{val}(\mathbf{A}^{(i-1)}, \mathbf{B}_1)$, $p_0 = p$ and $p_1 = 1 - p$, and plugging into Equation (28) yield that

$$\mathbb{E}_{\langle \mathbf{A}^{(k)}, \mathbf{B} \rangle} \left[\widehat{\mathcal{L}}_{\mathbf{A}}^{\Pi, \boldsymbol{\eta}} \right] \geq \frac{\sum_{j=0}^z \alpha_j \prod_{t=0}^{j-1} (1 - \alpha_t) \left(p \cdot \prod_{t=0}^{j-1} (1 - \beta_t^0) + (1 - p) \cdot \prod_{t=0}^{j-1} (1 - \beta_t^1) \right)^{k+1}}{\prod_{i=0}^{k-1} \left(p \cdot \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_0) + (1 - p) \cdot \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_1) \right)}. \quad (29)$$

We conclude the proof of this case by observing that for every $i \in (k - 1)$ it holds that $\text{val}(\mathbf{A}^{(i)}, \mathbf{B}) = p \cdot \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_0) + (1 - p) \cdot \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_1)$, and using the next claim (proven in Section 3.7.2), analogous to Claim 3.35.

Claim 3.40. *If \mathbf{B} controls $\text{root}(\Pi)$, it holds that*

$$\prod_{t=0}^{j-1} (1 - \beta_t) = p \cdot \prod_{t=0}^{j-1} (1 - \beta_t^0) + (1 - p) \cdot \prod_{t=0}^{j-1} (1 - \beta_t^1).$$

\mathbf{B} controls $\text{root}(\Pi)$ and $\text{val}(\Pi_0) > \text{val}(\Pi_1) = 0$. In this case, Claims 3.33 and 3.38 yield that $\alpha_j = 0$ for any $j \in (z^1)$. Hence, it suffices to prove that

$$\mathbb{E}_{\langle \mathbf{A}^{(k)}, \mathbf{B} \rangle} \left[\widehat{\mathcal{L}}_{\mathbf{A}}^{\Pi, \boldsymbol{\eta}} \right] \geq \frac{\sum_{j=z^1+1}^z \alpha_j \prod_{t=0}^{j-1} (1 - \beta_t)^{k+1} (1 - \alpha_t)}{\prod_{i=0}^{k-1} \text{val}(\mathbf{A}^{(i)}, \mathbf{B})}. \quad (30)$$

Thus, the proof immediately follows if $z^1 = z$, and in the following we assume that $z^1 < z$. As in Equation (27), compute

$$\begin{aligned} \mathbb{E}_{\langle \mathbf{A}^{(k)}, \mathbf{B} \rangle} \left[\widehat{\mathcal{L}}_{\mathbf{A}}^{\Pi, \boldsymbol{\eta}} \right] &= p \cdot \mathbb{E}_{\langle (\mathbf{A}^{(k)}, \mathbf{B})_0 \rangle} \left[\left(\widehat{\mathcal{L}}_{\mathbf{A}}^{\Pi, \boldsymbol{\eta}} \right)_0 \right] + (1 - p) \cdot \mathbb{E}_{\langle (\mathbf{A}^{(k)}, \mathbf{B})_1 \rangle} \left[\left(\widehat{\mathcal{L}}_{\mathbf{A}}^{\Pi, \boldsymbol{\eta}} \right)_1 \right] \\ &\geq p \cdot \frac{\sum_{j=0}^z \alpha_j^0 \prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1} (1 - \alpha_t^0)}{\prod_{i=0}^{k-1} \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_0)}, \end{aligned} \quad (31)$$

where the inequality follows Equation (22) and Claim 3.36. Claim 3.38 now yields

$$\mathbb{E}_{\langle \mathbf{A}^{(k)}, \mathbf{B} \rangle} \left[\widehat{\mathcal{L}}_{\mathbf{A}}^{\Pi, \boldsymbol{\eta}} \right] \geq \sum_{j=0}^z \alpha_j \prod_{t=0}^{j-1} (1 - \alpha_t) \cdot \frac{p \cdot \prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1}}{\prod_{i=0}^{k-1} \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_0)}, \quad (32)$$

where Claim 3.38 yields

$$\mathbb{E}_{\langle \mathbf{A}^{(k)}, \mathbf{B} \rangle} \left[\widehat{\mathcal{L}}_{\mathbf{A}}^{\Pi, \boldsymbol{\eta}} \right] \geq \sum_{j=z^1+1}^z \alpha_j \prod_{t=0}^{j-1} (1 - \alpha_t) \cdot \frac{p \cdot \prod_{t=0}^{j-1} (1 - \beta_t^0)^{k+1}}{\prod_{i=0}^{k-1} \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_0)}. \quad (33)$$

Multiplying both the numerator and the denominator for every summand of Equation (33) with p^k yields

$$\mathbb{E}_{\langle \mathbf{A}^{(k)}, \mathbf{B} \rangle} \left[\widehat{\mathcal{L}}_{\mathbf{A}}^{\Pi, \boldsymbol{\eta}} \right] \geq \sum_{j=z^1+1}^z \alpha_j \prod_{t=0}^{j-1} (1 - \alpha_t) \cdot \frac{\left(p \cdot \prod_{t=0}^{j-1} (1 - \beta_t^0) \right)^{k+1}}{\prod_{i=0}^{k-1} p \cdot \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_0)}. \quad (34)$$

Equation (30), and hence the proof of this case, is derived by observing that $\text{val}(\mathbf{A}^{(i)}, \mathbf{B}) = p \cdot \text{val}((\mathbf{A}^{(i)}, \mathbf{B})_0)$ for every $i \in (k-1)$,³⁵ and plugging Claims 3.39 and 3.40 into Equation (34).

B controls root(Π) and $\text{val}(\Pi_1) > \text{val}(\Pi_0) = 0$. Analogously to Claim 3.37, it holds that $\alpha_j^0 = 0$ for every $j \in (z)$. Claim 3.38 yields that $\alpha_j = 0$ for every $j \in (z)$. The proof of this case trivially follows since

$$\frac{\sum_{j=0}^z \alpha_j \prod_{t=0}^{j-1} (1 - \beta_t)^{k+1} (1 - \alpha_t)}{\prod_{i=0}^{k-1} \text{val}(\mathbf{A}^{(i)}, \mathbf{B})} = 0.$$

The above case analysis concludes the proof of the lemma when assuming that $e_\Pi(\lambda, b) \notin \{0, 1\}$ for both $b \in \{0, 1\}$. Assume that $e_\Pi(\lambda, b) = 1$ for some $b \in \{0, 1\}$. Since, by assumption, $\text{val}(\Pi) > 0$, it follows that $\text{val}(\Pi_b) > 0$. Moreover, the definition of conditional protocols (Definition 3.17) yields that $e_{\widehat{\Pi}_{(C,j)}^\eta}(\lambda, b) = 1$ and $e_{\widehat{\Pi}_{(C,j)}^\eta}(\lambda, 1 - b) = 0$ for any $(C, j) \in [(\mathbf{A}, z)]$ (regardless of which party controls $\text{root}(\Pi)$). By defining $\boldsymbol{\eta}_b = \boldsymbol{\eta}$, the definition of the dominated measure (Definition 3.10) yields that $\alpha_j = \alpha_j^b$ for every $j \in (z)$ and that $\beta_j = \beta_j^b$ for every $j \in (z-1)$. The proof of this case immediately follows from the induction hypothesis on Π_b . \square

3.7.2 Missing Proofs

This section is dedicated to proving deferred statements used in the proof of Lemma 3.28. We assume a fixed protocol Π , fixed real vector $\boldsymbol{\eta} = (\eta_{(\mathbf{A},0)}, \eta_{(\mathbf{B},0)}, \dots, \eta_{(\mathbf{B},z-1)}, \eta_{(\mathbf{A},z)})$ and a fixed positive integer k . We also assume that $\widehat{\Pi}_{(\mathbf{A},z)}^\eta \neq \perp$, $z^1 \leq z^0$ and $e_\Pi(\lambda, b) \in (0, 1)$ for both $b \in \{0, 1\}$. Recall that we defined two real vectors $\boldsymbol{\eta}_0$ and $\boldsymbol{\eta}_1$ (Definition 3.29), and for $b \in \{0, 1\}$ we defined $\alpha_j^b := \mu_{(\mathbf{A},j)}^{\Pi_b, \boldsymbol{\eta}_b}$ ($:= \mathbb{E}_{\langle \widehat{\Pi}_b \rangle_{(\mathbf{A},j)}^{\boldsymbol{\eta}_b}} [\widehat{M}_{(\mathbf{A},j)}^{\Pi_b, \boldsymbol{\eta}_b}]$) for $j \in (z)$, and $\beta_j^b := \mu_{(\mathbf{B},j)}^{\Pi_b, \boldsymbol{\eta}_b}$, for $j \in (z-1)$.

We begin with the following proposition, which underlies many of the claims to follow.

Proposition 3.41. *For $b \in \{0, 1\}$ and $(C, j) \in [(\mathbf{A}, z)]$, it holds that*

1. $(\widehat{\Pi}_{(C,j)}^\eta)_b = (\widehat{\Pi}_b)_{(C,j)}^{\boldsymbol{\eta}_b}$; and
2. $(\widehat{M}_{(C,j)}^{\Pi, \boldsymbol{\eta}})_b \equiv \widehat{M}_{(C,j)}^{\Pi_b, \boldsymbol{\eta}_b}$.

Namely, the restriction of $\widehat{\Pi}_{(C,j)}^\eta$ (the (C, j) 'th conditional protocol with respect to Π and $\boldsymbol{\eta}$) to its b 'th subtree is equal to the (C, j) 'th conditional protocol defined with respect to Π_b (b 'th subtree of Π) and $\boldsymbol{\eta}_b$. Moreover, the result of multiplying the C -dominated measure of $\widehat{\Pi}_{(C,j)}^\eta$ by $\eta_{(C,j)}$, and then restricting it to the subtree $(\widehat{\Pi}_{(C,j)}^\eta)_b$, is equivalent to multiplying the C -dominated measure of $(\widehat{\Pi}_b)_{(C,j)}^{\boldsymbol{\eta}_b}$ by $\eta_{(C,j)}^b$.³⁶

Proof of Proposition 3.41. The proof is by induction on the ordered pairs $[(\mathbf{A}, z)]$.

³⁵Recall that if $\text{val}(\mathbf{A}, \mathbf{B}) = 0$, then $\text{val}(\mathbf{A}^{(i)}, \mathbf{B}) = 0$ for every $i \in \mathbb{N}$.

³⁶Note that Item 1 is not immediate. Protocol $(\widehat{\Pi}_{(C,j)}^\eta)_b$ is a restriction of a protocol defined on the root of Π , whereas $(\widehat{\Pi}_b)_{(C,j)}^{\boldsymbol{\eta}_b}$ is a protocol defined on the root of Π_b .

Base case. Recall that the first pair of $[(A, z)]$ is $(A, 0)$. Definition 3.27 yields that $\widehat{\Pi}_{(A,0)}^\eta = \Pi$ and that $\left(\widehat{\Pi}_b\right)_{(A,0)}^{\eta_{\mathbf{b}}} = \Pi_b$, yielding that Item 1 holds for $(A, 0)$. As for Item 2, by Definition 3.10 and the assumption that $e_\Pi(\lambda, b) \in (0, 1)$ for both $b \in \{0, 1\}$, it holds that

$$\left(\widehat{M}_{(A,0)}^{\Pi, \eta}\right)_b \equiv \left(\eta_{(A,0)} \cdot M_{\Pi}^A\right)_b \equiv \begin{cases} \eta_{(A,0)} \cdot M_{\Pi_b}^A & \text{A controls } \text{root}(\Pi) \vee \text{Smaller}_\Pi(b); \\ \eta_{(A,0)} \cdot \frac{\xi_{(A,0)}^{1-b}}{\xi_{(A,0)}^b} \cdot M_{\Pi_b}^A & \text{otherwise.} \end{cases}$$

The proof that Item 2 holds for $(A, 0)$ now follows from Definition 3.29.

Induction step. Fix $(C, j) \in [(A, z)]$ and assume the claim holds for $\text{pred}(C, j)$. Using the induction hypothesis, we first prove Item 1 for (C, j) . Next, using the fact that Item 1 holds for (C, j) , we prove Item 2.

Proving Item 1. By Definition 3.27, it holds that

$$\begin{aligned} \left(\widehat{\Pi}_{(C,j)}^\eta\right)_b &= \left(\widehat{\Pi}_{\text{pred}(C,j)}^\eta \upharpoonright_{\neg} \left(\widehat{M}_{\text{pred}(C,j)}^{\Pi, \eta}\right)_b\right)_b \\ &= \left(\widehat{\Pi}_{\text{pred}(C,j)}^\eta\right)_b \upharpoonright_{\neg} \left(\widehat{M}_{\text{pred}(C,j)}^{\Pi, \eta}\right)_b \\ &= \left(\widehat{\Pi}_b\right)_{\text{pred}(C,j)}^{\eta_{\mathbf{b}}} \upharpoonright_{\neg} \left(\widehat{M}_{\text{pred}(C,j)}^{\Pi_b, \eta_{\mathbf{b}}}\right)_b \\ &= \left(\widehat{\Pi}_b\right)_{(C,j)}^{\eta_{\mathbf{b}}}, \end{aligned}$$

where the third equality follows from the induction hypothesis.

Proving Item 2. Similarly to the base case, Definition 3.10 yields that

$$\left(\widehat{M}_{(C,j)}^{\Pi, \eta}\right)_b \equiv \begin{cases} 0 & e_{\widehat{\Pi}_{(C,j)}^\eta}(\lambda, b) = 0; \\ \eta_{(C,j)} \cdot M_{\left(\widehat{\Pi}_{(C,j)}^\eta\right)_b}^C & e_{\widehat{\Pi}_{(C,j)}^\eta}(\lambda, b) = 1; \\ \eta_{(C,j)} \cdot M_{\left(\widehat{\Pi}_{(C,j)}^\eta\right)_b}^C & e_{\widehat{\Pi}_{(C,j)}^\eta}(\lambda, b) \notin \{0, 1\} \wedge \\ & \left(C \text{ controls } \text{root}(\Pi) \vee \text{Smaller}_{\widehat{\Pi}_{(C,j)}^\eta}(b)\right); \\ \eta_{(C,j)} \cdot \frac{\xi_{(C,j)}^{1-b}}{\xi_{(C,j)}^b} \cdot M_{\left(\widehat{\Pi}_{(C,j)}^\eta\right)_b}^C & \text{otherwise,} \end{cases}$$

and the proof follows by Item 1 and Definition 3.29. □

Recall that the real numbers α_j^b and β_j^b were defined to be the expected values of the (A, j) 'th and (B, j) 'th dominated measures in the sequence $(A, z, \eta_{\mathbf{b}})$ -DMS (Π_b) , respectively (see the proof of Lemma 3.28). Following Proposition 3.41, we could equivalently define α_j^b and β_j^b with respect to the sequence (A, z, η) -DMS (Π) .

Proposition 3.42. *For both $b \in \{0, 1\}$, it holds that*

1. $\alpha_j^b = \mathbb{E} \langle (\widehat{\Pi}_{(\mathbf{A},j)}^\eta)_b \rangle \left[\left(\widehat{M}_{(\mathbf{A},j)}^{\Pi, \eta} \right)_b \right]$ for every $j \in (z)$; and

2. $\beta_j^b = \mathbb{E} \langle (\widehat{\Pi}_{(\mathbf{B},j)}^\eta)_b \rangle \left[\left(\widehat{M}_{(\mathbf{B},j)}^{\Pi, \eta} \right)_b \right]$ for every $j \in (z-1)$.

Proof. Immediately follows Proposition 3.41. □

Proposition 3.42 allows us to use Proposition 3.13 in order to analyze the connections between α_j^0 and α_j^1 to α_j , and similarly between β_j^0 and β_j^1 to β_j . Towards this goal, we analyze the edge distribution of the conditional protocols defined in the procedure that generates the measure sequence (\mathbf{A}, z, η) -DMS (Π) .

Proposition 3.43. *The following holds for both $b \in \{0, 1\}$.*

1. **A controls** $\text{root}(\Pi) \implies$

$$(a) \ e_{\widehat{\Pi}_{(\mathbf{A},j)}^\eta}(\lambda, b) = e_\Pi(\lambda, b) \cdot \frac{\prod_{t=0}^{j-1} (1 - \alpha_t^b)}{\prod_{t=0}^{j-1} (1 - \alpha_t)}$$

$$(b) \ e_{\widehat{\Pi}_{(\mathbf{B},j)}^\eta}(\lambda, b) = e_\Pi(\lambda, b) \cdot \frac{\prod_{t=0}^j (1 - \alpha_t^b)}{\prod_{t=0}^j (1 - \alpha_t)}$$

2. **B controls** $\text{root}(\Pi) \implies$

$$(a) \ e_{\widehat{\Pi}_{(\mathbf{A},j)}^\eta}(\lambda, b) = e_\Pi(\lambda, b) \cdot \frac{\prod_{t=0}^{j-1} (1 - \beta_t^b)}{\prod_{t=0}^{j-1} (1 - \beta_t)}$$

$$(b) \ e_{\widehat{\Pi}_{(\mathbf{B},j)}^\eta}(\lambda, b) = e_\Pi(\lambda, b) \cdot \frac{\prod_{t=0}^{j-1} (1 - \beta_t^b)}{\prod_{t=0}^{j-1} (1 - \beta_t)}$$

Proof. We prove Item 1 using induction on the ordered pairs $[(\mathbf{A}, z)]$. The proof of Item 2 is analogous.

Base case. The proof follows since according to Definition 3.27, it holds that $\widehat{\Pi}_{(\mathbf{A},0)}^\eta = \Pi$.

Induction step. Fix $(\mathbf{C}, j) \in [(\mathbf{A}, z)]$ and assume the claim holds for $\text{pred}(\mathbf{C}, j)$. The proof splits according to which party \mathbf{C} is.

Case $\mathbf{C} = \mathbf{A}$. If $e_{\widehat{\Pi}_{(\mathbf{B},j-1)}^\eta}(\lambda, b) = 0$, Definition 3.17 yields that $e_{\widehat{\Pi}_{(\mathbf{A},j)}^\eta}(\lambda, b) = 0$. The proof follows since, by the induction hypothesis, it holds that

$$e_{\widehat{\Pi}_{(\mathbf{A},j)}^\eta}(\lambda, b) = e_{\widehat{\Pi}_{(\mathbf{B},j-1)}^\eta}(\lambda, b) = e_\Pi(\lambda, b) \cdot \frac{\prod_{t=0}^{j-1} (1 - \alpha_t^b)}{\prod_{t=0}^{j-1} (1 - \alpha_t)}.$$

In the complementary case, i.e., $e_{\widehat{\Pi}_{(\mathbf{B},j-1)}^\eta}(\lambda, b) > 0$, Proposition 3.13 and Definition 3.10 yield that $\beta_{j-1} = \beta_{j-1}^b$. It must be the case that $\beta_{j-1} = \beta_{j-1}^b < 1$, since otherwise, according to

Definition 3.27, it holds that $\widehat{\Pi}_{(A,j)}^\eta = \perp$, a contradiction to the assumption that $\widehat{\Pi}_{(A,z)}^\eta \neq \perp$. The proof follows since in this case Definition 3.17 and Proposition 3.42 yield that

$$\begin{aligned} e_{\widehat{\Pi}_{(A,j)}^\eta}(\lambda, b) &= e_{\widehat{\Pi}_{(B,j-1)}^\eta}(\lambda, b) \cdot \frac{1 - \beta_{j-1}^b}{1 - \beta_{j-1}} \\ &= e_{\widehat{\Pi}_{(B,j-1)}^\eta}(\lambda, b) \\ &= e_{\Pi}(\lambda, b) \cdot \frac{\prod_{t=0}^{j-1} (1 - \alpha_t^b)}{\prod_{t=0}^{j-1} (1 - \alpha_t)}, \end{aligned}$$

where the last equality follows the induction hypothesis.

Case C = B. It must be that case that $\alpha_j < 1$, since otherwise, similarly to the previous case and according to Definition 3.27, it holds that $\widehat{\Pi}_{(B,j)}^\eta = \perp$, a contradiction to the assumption that $\widehat{\Pi}_{(A,z)}^\eta \neq \perp$. The proof follows since in this case Definition 3.17 and Proposition 3.42 yield that

$$\begin{aligned} e_{\widehat{\Pi}_{(B,j)}^\eta}(\lambda, b) &= e_{\widehat{\Pi}_{(A,j)}^\eta}(\lambda, b) \cdot \frac{1 - \alpha_j^b}{1 - \alpha_j} \\ &= e_{\Pi}(\lambda, b) \cdot \frac{\prod_{t=0}^{j-1} (1 - \alpha_t^b)}{\prod_{t=0}^{j-1} (1 - \alpha_t)} \cdot \frac{1 - \alpha_j^b}{1 - \alpha_j} \\ &= e_{\Pi}(\lambda, b) \cdot \frac{\prod_{t=0}^j (1 - \alpha_t^b)}{\prod_{t=0}^j (1 - \alpha_t)}, \end{aligned}$$

where the second equality follows the induction hypothesis. □

Using the above propositions, we now turn our focus to proving the claims in the proof of Lemma 3.28. To facilitate reading and tracking the proof, we cluster claims together according to their role in the proof of Lemma 3.28.

3.7.2.1 Proving Claims 3.30 and 3.32

Proof of Claim 3.30. For $b \in \{0, 1\}$ it holds that

$$\begin{aligned} \widehat{L}_A^{\Pi_b, \eta_b} &\equiv \sum_{j=0}^z \widehat{M}_{(A,j)}^{\Pi_b, \eta_b} \cdot \prod_{t=0}^{j-1} \left(1 - \widehat{M}_{(A,t)}^{\Pi_b, \eta_b}\right) \\ &\equiv \sum_{j=0}^z \left(\widehat{M}_{(A,j)}^{\Pi, \eta}\right)_b \cdot \prod_{t=0}^{j-1} \left(1 - \left(\widehat{M}_{(A,t)}^{\Pi, \eta}\right)_b\right) \\ &\equiv \left(\widehat{L}_A^{\Pi, \eta}\right)_b, \end{aligned}$$

where the second equality follows Proposition 3.41. □

Proof of Claim 3.32. Assume towards a contradiction that $z^0 < z$. By the definition of z^0 (Definition 3.31) and the definition of conditional protocols (Definition 3.17), it follows that

$\left(\widehat{\Pi}_0\right)_{(A, z^0+1)}^{\eta_0} = \perp$. Since (by assumption) $z^1 \leq z^0$, it also holds that $\left(\widehat{\Pi}_1\right)_{(A, z^0+1)}^{\eta_1} = \perp$. Hence, Proposition 3.41 yields that $\left(\widehat{\Pi}_{(A, z^0+1)}^{\eta}\right)_0, \left(\widehat{\Pi}_{(A, z^0+1)}^{\eta}\right)_1 = \perp$. Namely, the function describing $\widehat{\Pi}_{(A, z^0+1)}^{\eta}$ does not correspond to any two-party execution when restricting it to the subtrees $\mathcal{T}(\Pi_0)$ and $\mathcal{T}(\Pi_1)$. Hence, the aforementioned function does not correspond to a two-party execution (over $\mathcal{T}(\Pi)$), in contradiction to the assumption that $\widehat{\Pi}_{(A, z)}^{\eta} \neq \perp$. \square

3.7.2.2 Proving Claims 3.33 to 3.35

The following proofs rely on the next observation. As long as $\alpha_j^b < 1$ and $\beta_j^b < 1$, Proposition 3.43 ensures that there is a positive probability to visit both the left and the right subtree of the (C, j) 'th conditional protocol.

Proof of Claim 3.34. Assume that A controls $\text{root}(\Pi)$ and that $z^1 < z$. Assume towards a contradiction that $\alpha_{z^1}^1 < 1$. Since $z^1 \leq z^0$ (by assumption), it follows that $\alpha_{z^1}^0 < 1$ as well. The definition of z^1 (Definition 3.31) yields that $\beta_{z^1}^1 = 1$. However, Proposition 3.43 yields that $e_{\widehat{\Pi}_{(B, j)}^{\eta}}(\lambda, b) \in (0, 1)$ for both $b \in \{0, 1\}$, and thus Propositions 3.13 and 3.42 yield that $\beta_{z^1} = 1$. Now, Definition 3.27 yields that $\widehat{\Pi}_{(A, z^1+1)}^{\eta} = \perp$, a contradiction to the assumption that $\widehat{\Pi}_{(A, z)}^{\eta} \neq \perp$. \square

Proof of Claim 3.33. For $j \in (z^1 - 1)$, it holds that $e_{\widehat{\Pi}_{(B, j)}^{\eta}}(\lambda, b) \in (0, 1)$ for both $b \in \{0, 1\}$. Thus, $\beta_j^0 = \beta_j^1 = \beta_j$ is a direct implication of Propositions 3.13 and 3.41.

For $z^1 \leq z - 1$, Claim 3.34 and Proposition 3.43 yield that $e_{\widehat{\Pi}_{(B, j)}^{\eta}}(\lambda, 0) = 1$. Since, by Definition 3.29, it holds that $\eta_{(B, j)} = \eta_{(B, j)}^0$, Definition 3.10 and Proposition 3.41 yield that $\beta_j^0 = \beta_j$. \square

Proof of Claim 3.35. The proof immediately follows from Propositions 3.42 and 3.43. \square

3.7.2.3 Proving Claims 3.36 and 3.37

Proof of Claim 3.36. By Definition 3.10 it holds that $\widehat{M}_{(A, j)}^{\Pi_1, \eta_1} \equiv 0$ for every $j \in (z)$. Definition 3.27 yields that $\widehat{L}_A^{\Pi_1, \eta_1} \equiv 0$. The proof follows from Claim 3.30. \square

Proof of Claim 3.37. Follows similar arguments to the above proof of Claim 3.36, together with Proposition 3.42. \square

3.7.2.4 Proving Claims 3.38 to 3.40

The proofs of the rest of the claims stated in the proof of Lemma 3.28 are analogous to the claims proven above. Specifically, Claim 3.38 is analogous to Claim 3.33, Claim 3.39 is analogous to Claim 3.34, and Claim 3.40 is analogous to Claim 3.35.

3.8 Proving Lemma 3.26

Lemma 3.26 immediately follows by the next lemma.

Lemma 3.44. *For every protocol Π , there exists $(C, j) \in \{\mathbf{A}, \mathbf{B}\} \times \mathbb{N}$ such that*

$$\mathbb{E}_{\langle \Pi_{(C,j)} \rangle} \left[M_{\Pi_{(C,j)}}^C \right] = 1.$$

The proof of Lemma 3.44 is given below, but first we use it to derive Lemma 3.26.

Proof of Lemma 3.26. Let z be the minimal integer such that $\sum_{j=0}^z \alpha_j \geq c$ or $\sum_{j=0}^z \beta_j \geq c$. Note that such z is guaranteed to exist by Lemma 3.44 and since by Lemma 3.11 it holds that $\alpha_j = \mathbb{E}_{\langle \Pi_{(\mathbf{A},j)} \rangle} \left[M_{\Pi_{(\mathbf{A},j)}}^{\mathbf{A}} \right]$ and $\beta_j = \mathbb{E}_{\langle \Pi_{(\mathbf{B},j)} \rangle} \left[M_{\Pi_{(\mathbf{B},j)}}^{\mathbf{B}} \right]$. The proof splits to the following cases.

Case $\sum_{j=0}^z \alpha_j \geq c$. By the choice of z it holds that $\sum_{j=0}^{z-1} \alpha_j < c$ and $\sum_{j=0}^{z-1} \beta_j < c$. Lemma 3.23 yields that

$$\begin{aligned} \mathbb{E}_{\langle \Pi \rangle} \left[L_{\Pi}^{\mathbf{A},z} \right] &= \sum_{j=0}^z \alpha_j \prod_{t=0}^{j-1} (1 - \beta_t)(1 - \alpha_t) \\ &\geq \left(\sum_{j=0}^z \alpha_j \right) \cdot \left(1 - \sum_{j=0}^{z-1} \beta_j \right) \cdot \left(1 - \sum_{j=0}^{z-1} \alpha_j \right) \\ &\geq c \cdot (1 - 2c), \end{aligned}$$

where the first inequality follows by multiplying the j 'th summand by $\prod_{t=j}^{z-1} (1 - \beta_t)(1 - \alpha_t) \leq 1$ and both inequalities follow since $(1 - x)(1 - y) \geq 1 - (x + y)$ for any $x, y \geq 0$. Hence, z satisfies Item 1.

Case $\sum_{j=0}^z \alpha_j < c$. By the choice of z it holds that $\sum_{j=0}^z \beta_j \geq c$ and $\sum_{j=0}^{z-1} \beta_j < c$. Similar arguments to the previous case show that z satisfies Item 2. □

Towards proving Lemma 3.44 we prove that there is always a leaf for which the value of the dominated measure is 1.

Claim 3.45. *Let Π be a protocol with $\text{OPT}_{\mathbf{A}}(\Pi) = 1$. Then there exists $\ell \in \mathcal{L}_1(\Pi)$ such that $M_{\Pi}^{\mathbf{A}}(\ell) = 1$.*

Proof. The proof is by induction on the round complexity of Π .

Assume that $\text{round}(\Pi) = 0$ and let ℓ be the only node in $\mathcal{T}(\Pi)$. Since $\text{OPT}_{\mathbf{A}}(\Pi) > 0$, it must be the case that $\chi_{\Pi}(\ell) = 1$. The proof follows since Definition 3.10 yields that $M_{\Pi}^{\mathbf{A}}(\ell) = 1$.

Assume that $\text{round}(\Pi) = m + 1$ and that the lemma holds for m -round protocols. If $e_{\Pi}(\lambda, b) = 1$ for some $b \in \{0, 1\}$, then by Proposition 3.8 it holds that $\text{OPT}_{\mathbf{A}}(\Pi_b) = \text{OPT}_{\mathbf{A}}(\Pi) = 1$. This allows us to apply the induction hypothesis on Π_b , which yields that there exists $\ell \in \mathcal{L}_1(\Pi_b)$ such that $M_{\Pi_b}^{\mathbf{A}}(\ell) = 1$. In this case, according to Definition 3.10, $M_{\Pi}^{\mathbf{A}}(\ell) = M_{\Pi_b}^{\mathbf{A}}(\ell) = 1$, and the proof follows.

In the following we assume that $e_{\Pi}(\lambda, b) \in (0, 1)$ for any $b \in \{0, 1\}$. We conclude the proof using the following case analysis.

A controls $\text{root}(\Pi)$. According to Proposition 3.8, there exists $b \in \{0, 1\}$ such that $\text{OPT}_{\mathbf{A}}(\Pi_b) = \text{OPT}_{\mathbf{A}}(\Pi) = 1$. This allows us to apply the induction hypothesis on Π_b , which yields that there exists $\ell \in \mathcal{L}_1(\Pi_b)$ such that $M_{\Pi_b}^{\mathbf{A}}(\ell) = 1$. The \mathbf{A} -maximal property of $M_{\Pi}^{\mathbf{A}}$ (Proposition 3.13(1)) yields that $M_{\Pi}^{\mathbf{A}}(\ell) = M_{\Pi_b}^{\mathbf{A}}(\ell) = 1$, and the proof for this case follows.

B controls $\text{root}(\Pi)$. According to Proposition 3.8, $\text{OPT}_A(\Pi_b) = \text{OPT}_A(\Pi) = 1$ for both $b \in \{0, 1\}$.

This allows us to apply the induction hypothesis on Π_0 and Π_1 , which yields that there exists $\ell_0 \in \mathcal{L}_1(\Pi_0)$ and $\ell_1 \in \mathcal{L}_1(\Pi_1)$ such that $M_{\Pi_0}^A(\ell_0) = 1$ and $M_{\Pi_1}^A(\ell_1) = 1$. The B-minimal property of M_{Π}^A (Proposition 3.13(2)) yields that there exists $b \in \{0, 1\}$ such that $M_{\Pi}^A(\ell_b) = M_{\Pi_b}^A(\ell_b) = 1$ (the bit b for which $\text{Smaller}_{\Pi}(b) = 1$), and the proof for this case follows.

This concludes the case analysis and the proof follows. \square

We can now derive Lemma 3.44. Claim 3.45 and Proposition 3.13 yield that the number of possible transcripts of $\Pi_{(C,j)}$ shrinks as (C, j) grows. Specifically, at least one possible transcript of $\Pi_{(A,j)}$ whose common outcome is 1 (the transcript represented by the leaf is guaranteed to exist from Claim 3.45) is *not* a possible transcript of $\Pi_{(B,j)}$. Similarly, at least one possible transcript of $\Pi_{(B,j-1)}$ whose common outcome is 0 is not a possible transcript of $\Pi_{(A,j)}$. Since the number of possible transcripts of Π is finite (though might be exponentially large), there exists $j \in \mathbb{N}$ such that either the common outcome of all possible transcripts $\Pi_{(A,j)}$ is 1 or the common outcome of all possible transcripts of $\Pi_{(B,j)}$ is 0. The expected value of the A-dominated measure of $\Pi_{(A,j)}$ or the B-dominated measure of $\Pi_{(B,j)}$ will be 1. The formal proof is given next.

Proof of Lemma 3.44. Assume towards a contradiction that $\mathbb{E}_{\langle \Pi_{(C,j)} \rangle} \left[M_{\Pi_{(C,j)}}^C \right] < 1$ for every $(C, j) \in \{A, B\} \times \mathbb{N}$. It follows that $\Pi_{(C,j)} \neq \perp$ for every such (C, j) . For a pair $(C, j) \in \{A, B\} \times \mathbb{N}$, recursively define $\mathcal{L}_{(C,j)} := \mathcal{L}_{\text{pred}(C,j)} \cup \mathcal{S}_{(C,j)}$, where $\mathcal{S}_{(C,j)} := \left\{ \ell \in \mathcal{L}(\Pi) : M_{\Pi_{(C,j)}}^C(\ell) = 1 \right\}$ and $\mathcal{L}_{(B,-1)} := \emptyset$. The following claim (proven below) shows two properties of $\mathcal{S}_{(C,j)}$.

Claim 3.46. *It holds that $\mathcal{S}_{(C,j)} \neq \emptyset$ and $\mathcal{L}_{\text{pred}(C,j)} \cap \mathcal{S}_{(C,j)} = \emptyset$ for every $(C, j) \succeq (B, 0)$.*

Claim 3.46 yields that $|\mathcal{L}_{(C,j)}| > |\mathcal{L}_{\text{pred}(C,j)}|$ for every $(C, j) \succeq (B, 0)$, a contradiction to the fact that $\mathcal{L}_{(C,j)} \subseteq \mathcal{L}(\Pi)$ for every (C, j) . \square

Proof of Claim 3.46. Let $(C, j) \succeq (B, 0)$. By Lemma 3.20 it holds that $\text{OPT}_C(\Pi_{(C,j)}) = 1$.³⁷ Hence, Claim 3.45 yields that $\mathcal{S}_{(C,j)} \neq \emptyset$.

Towards proving the second property, let $\ell' \in \mathcal{L}_{\text{pred}(C,j)}$, and let $(C', j') \in [\text{pred}(C, j)]$ such that $\ell' \in \mathcal{S}_{(C',j')}$. By the definition of $\mathcal{S}_{(C',j')}$, it holds that $M_{\Pi_{(C',j')}}^C(\ell') = 1$. By Proposition 3.18 it holds that $\ell' \notin \text{Supp}(\langle \Pi_{(C'',j'')} \rangle)$ for every $(C'', j'') \succ (C', j')$. Since $(C, j) \succ \text{pred}(C, j) \succeq (C', j')$, it holds that $\ell' \notin \text{Supp}(\langle \Pi_{(C,j)} \rangle)$. By Definition 3.10 it holds that $M_{\Pi_{(C,j)}}^C(\ell) = 0$ for every $\ell \notin \text{Supp}(\langle \Pi_{(C,j)} \rangle)$, and thus $\ell' \notin \mathcal{S}_{(C,j)}$. Hence, $\mathcal{L}_{\text{pred}(C,j)} \cap \mathcal{S}_{(C,j)} = \emptyset$. \square

4 Efficiently Biasing Coin-Flipping Protocols

In Section 3 we showed that for any coin-flipping protocol and any $\varepsilon \in (0, \frac{1}{2}]$, applying $\kappa = \kappa(\varepsilon)$ recursions of the biased-continuation attack biases the honest party's outcome by (at least) $1/2 - \varepsilon$.

³⁷Note that this might not hold for $\Pi_{(A,0)} = \Pi$. Namely, it might be the case that $\text{OPT}_B(\Pi) = 1$. In this case M_{Π}^A is the zero measure, $\Pi_{(B,0)} = \Pi$ and $\mathcal{S}_{(A,0)} = \emptyset$.

Implementing this attack, however, requires access to a sampling algorithm (i.e., `BiasedCont`; see Definition 3.1), which we do not know how to efficiently implement even when assuming OWFs do not exist. In this section we show that the nonexistence of OWFs does suffice to implement an approximation of `BiasedCont`, which in turn can be used to implement a strong enough variant of the aforementioned attack.

The outline of this section is as follows. We begin, in Section 4.1, by defining an approximation of the `BiasedCont` sampling algorithm that can be efficiently implemented assuming the nonexistence of OWFs. We then use this approximation to define the recursive approximated biased-continuation attacker, the approximated variant of the recursive biased-continuation attacker defined in Section 3. We then relate the success probability of this attacker to the probability that it visits low-value nodes (the expected protocol's outcome conditioned on visiting the nodes [transcripts] is close to zero), and to the probability that it visits unbalanced nodes (the attack drastically increases the probability of visiting these nodes). Finally, we relate these two probabilities to one another by showing that it is unlikely that a protocol will visit an unbalanced node without first visiting a low-value node. We conclude that the recursive approximated biased-continuation attacker successfully biases protocols that have no low-value nodes. In Section 4.2 we define a special class of protocols, called *pruned* protocols, that have (almost) no low-value nodes. We use the observations made in Section 4.1 to prove that the recursive approximated biased-continuation attacker performs well on such pruned protocols. In Section 4.3 we define the *pruning-in-the-head attacker*, which behaves as if the protocol it is attacking is pruned, and by doing so manages to make use of the recursive approximated biased-continuation attacker to attack *any* protocol. Finally, in Section 4.4 we show that the assumption that OWFs do not exist indeed implies that the above attacker can be implemented efficiently, and thus that the outcome on any coin-flipping protocol can be biased to be arbitrarily close to 0 or 1.

4.1 The Approximated Biased-Continuation Attacker

We require the approximated biased-continuator sampler to work well only when applied on nodes whose value is not too close to the borders. (This value is the probability that the protocol outcome is 1 given that the current transcript is the node's label.) In the following let BiasedCont_Π be as in Definition 3.1.

Definition 4.1 ($\text{BiasedCont}_\Pi^{\xi, \delta}$). *Algorithm BC is a (ξ, δ) -biased-continuator for an m -round protocol Π if the following hold.*

1. $\Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1) : \text{SD}(\text{BC}(\ell_{1, \dots, i}, 1), \text{BiasedCont}_\Pi(\ell_{1, \dots, i}, 1)) > \xi \wedge \text{val}(\Pi_{\ell_{1, \dots, i}}) > \delta] \leq \xi$, and
2. $\Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1) : \text{SD}(\text{BC}(\ell_{1, \dots, i}, 0), \text{BiasedCont}_\Pi(\ell_{1, \dots, i}, 0)) > \xi \wedge \text{val}(\Pi_{\ell_{1, \dots, i}}) < 1 - \delta] \leq \xi$.

Let $\text{BiasedCont}_\Pi^{\xi, \delta}$ be an arbitrary (but fixed) (ξ, δ) -biased-continuator of Π .

The recursive approximated biased-continuation attacker is identical to that defined in Section 3, except that it uses the approximated biased-continuator sampler and not the ideal one.

Let $A_\Pi^{(0, \xi, \delta)} \equiv A$, and for integer $i > 0$ define:

Algorithm 4.2 ($A_\Pi^{(i, \xi, \delta)}$).

Input: transcript $u \in \{0, 1\}^$.*

Operation:

1. If $u \in \mathcal{L}(\Pi)$, output $\chi_\Pi(u)$ and halt.
2. Set $\text{msg} = \text{BiasedCont}_{\left(\mathbf{A}_\Pi^{(i-1, \xi, \delta)}, \mathbf{B}\right)}^{\xi, \delta}(u, 1)$.
3. Send msg to \mathbf{B} .
4. If $u' = u \circ \text{msg} \in \mathcal{L}(\Pi)$, output $\chi_\Pi(u')$.

Adversary $\mathbf{B}_\Pi^{(i, \xi, \delta)}$ attacking towards zero is analogously defined. In the following we sometimes refer to the base (non-recursive) version of the above algorithm, i.e., $\mathbf{A}_\Pi^{(1, \xi, \delta)}$, as the approximated biased-continuation attacker. As in Section 3, when clear from the context, we remove the protocol name (i.e., Π) from the subscript of the above attacker.³⁸

Our first goal is to bound the difference between the (non-recursive) biased-continuation attacker and its approximated variant defined above. Clearly, if the statistical distance of the answers of BiasedCont and $\text{BiasedCont}^{\xi, \delta}$ is small, then so will be the difference between the attackers. Definition 4.1, however, does not always guarantee such small statistical distance. Specifically, there is no such guarantee for low-value and high-value transcripts.

Definition 4.3 (low-value and high-value nodes). *For a protocol $\Pi = (\mathbf{A}, \mathbf{B})$ and $\delta \in [0, 1]$, let*

- $\text{Small}_\Pi^\delta = \{u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi) : \text{val}(\Pi_u) \leq \delta\}$, and
- $\text{Large}_\Pi^\delta = \{u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi) : \text{val}(\Pi_u) \geq 1 - \delta\}$.

For $\mathbf{C} \in \{\mathbf{A}, \mathbf{B}\}$, let $\text{Small}_\Pi^{\delta, \mathbf{C}} = \text{Small}_\Pi^\delta \cap \text{Ctrl}_\Pi^\mathbf{C}$ and similarly let $\text{Large}_\Pi^{\delta, \mathbf{C}} = \text{Large}_\Pi^\delta \cap \text{Ctrl}_\Pi^\mathbf{C}$.³⁹

For non-low-value and non-high-value transcripts, Definition 4.1 guarantees small statistical distance between the answers of BiasedCont and $\text{BiasedCont}^{\xi, \delta}$, when queried on transcripts chosen according to the honest distribution of leaves (i.e., $\langle \Pi \rangle$). The queries the biased-continuation attacker makes, however, might be chosen from a different distribution, making some transcripts much more likely to be queried than before. We call such transcripts “unbalanced”.

Definition 4.4 (unbalanced nodes). *For a protocol $\Pi = (\mathbf{A}, \mathbf{B})$ and $\gamma \geq 1$, let $\text{UnBal}_\Pi^\gamma = \left\{u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi) : \mathbf{v}_{\left(\mathbf{A}_\Pi^{(1)}, \mathbf{B}\right)}(u) \geq \gamma \cdot \mathbf{v}_{(\mathbf{A}, \mathbf{B})}(u)\right\}$, where $\mathbf{A}_\Pi^{(1)}$ is as in Algorithm 3.2 and \mathbf{v} as in Definition 2.2.*

Namely, UnBal_Π^γ are those nodes that a random execution of $(\mathbf{A}^{(1)}, \mathbf{B})$ visits with probability at least γ times the probability that a random execution of Π does.

Consider an execution of $(\mathbf{A}^{(1, \xi, \delta)}, \mathbf{B})$. Such an execution asks $\text{BiasedCont}^{\xi, \delta}$ for continuations of transcripts under \mathbf{A} 's control, leading to 1-leaves. Hence, as long as this execution generates neither low-value transcripts under \mathbf{A} 's control nor unbalanced transcripts, we expect the approximated biased-continuation attacker to do almost as well as its ideal variant. This is formally put in the following lemma.

³⁸As a rule of thumb, in statements and definitions we explicitly write the protocols to which the algorithms refer, whereas in proofs and informal discussions we usually omit them.

³⁹Recall that $\text{Ctrl}_\Pi^\mathbf{C}$ denotes the nodes in $\mathcal{T}(\Pi)$ controlled by party \mathbf{C} (see Definition 2.2).

Lemma 4.5. *Let $\Pi = (A, B)$ be an m -round protocol and let $\delta \in (0, \frac{1}{2}]$. Then*

$$\text{SD} \left(\langle A_{\Pi}^{(1)}, B \rangle, \langle A_{\Pi}^{(1, \xi, \delta)}, B \rangle \right) \leq m \cdot \gamma \cdot \left(2\xi + \Pr_{\langle A, B \rangle} \left[\text{desc}(\text{Small}_{\Pi}^{\delta, A}) \right] \right) + \Pr_{\langle A_{\Pi}^{(1)}, B \rangle} \left[\text{desc}(\text{UnBal}_{\Pi}^{\gamma}) \right] \quad 40$$

for every $\gamma \geq 1$ and $\xi > 0$.

Proof. We use Lemma 2.14. For function O , let H^O be an algorithm that outputs the transcript of a random execution of $(A_{\Pi}^{(1)}, B)$ in which $A_{\Pi}^{(1)}$'s calls to BiasedCont_{Π} are sent to O instead. Let f and g be the (random) functions BiasedCont_{Π} and $\text{BiasedCont}_{\Pi}^{\xi, \delta}$ respectively, with the exception that $f(\perp) = g(\perp) = \perp$. By construction, it holds that

$$\text{SD} \left(\langle A_{\Pi}^{(1)}, B \rangle, \langle A_{\Pi}^{(1, \xi, \delta)}, B \rangle \right) = \text{SD} \left(H^f, H^g \right). \quad (35)$$

For $i \in [m]$, let D'_i be the distribution of the i 'th node under A 's control in a random execution of Π , taking the value \perp if no such node exists, and let $D_i = (D'_i, 1)$, with $(\perp, 1) = \perp$. By definition,

$$\begin{aligned} \mathbb{E}_{d \leftarrow D_i} [\text{SD}(f(d), g(d))] &= \mathbb{E}_{d \leftarrow D_i} \left[\text{SD}(\text{BiasedCont}_{\Pi}(d), \text{BiasedCont}_{\Pi}^{\xi, \delta}(d)) \cdot 1_{\neg \perp}(d) \right] \\ &\leq 2\xi + \Pr_{\langle \Pi \rangle} \left[\text{desc}(\text{Small}_{\Pi}^{\delta, A}) \right], \end{aligned} \quad (36)$$

letting the indicator $1_{\neg \perp}(d)$ take the value one if $d \neq \perp$, and zero otherwise.

Let Q_i denote the i 'th query to f in a random execution of H^f , taking the value \perp if no such query exists, and let $Q = (Q_1, \dots, Q_m)$. By definition,

$$\Pr_{(q_1, \dots, q_m) \leftarrow Q} [\exists i \in [m]: q_i \neq \perp \wedge Q_i(q_i) > \gamma \cdot D_i(q_i)] = \Pr_{\langle A_{\Pi}^{(1)}, B \rangle} \left[\text{desc}(\text{UnBal}_{\Pi}^{\gamma}) \right]. \quad (37)$$

Hence, the proof follows by Lemma 2.14, letting $k := m$, $a := 2\xi + \Pr_{\langle \Pi \rangle} \left[\text{desc}(\text{Small}_{\Pi}^{\delta, A}) \right]$, $\lambda := \gamma$ and $b := \Pr_{\langle A_{\Pi}^{(1)}, B \rangle} \left[\text{desc}(\text{UnBal}_{\Pi}^{\gamma}) \right]$. \square

In the rest of this subsection we show that it is unlikely that a protocol will visit an unbalanced node without first visiting a low-value node controlled by the attacking party. This fact is used in Sections 4.2 and 4.3 for designing an effective attack using the approximated biased-continuation attacker.

4.1.1 Bounding the Probability of Visiting Unbalanced Nodes

Given a protocol $\Pi = (A, B)$, we would like to understand what makes a node unbalanced. Let u be a γ -unbalanced node, i.e., $v_{(A^{(1)}, B)}(u) \geq \gamma \cdot v_{(A, B)}(u)$. By the edge distribution of $(A^{(1)}, B)$ (Claim 3.4), it follows that

$$\frac{v_{(A^{(1)}, B)}(u)}{v_{(A, B)}(u)} = \prod_{\substack{0 \leq i \leq |u| - 1: \\ u_1, \dots, i \in \text{Ctrl}_{\Pi}^A}} \frac{\text{val}(\Pi_{u_1, \dots, i+1})}{\text{val}(\Pi_{u_1, \dots, i})} \geq \gamma. \quad (38)$$

⁴⁰Recall that for $S \subseteq \mathcal{V}(\Pi)$, $\text{desc}(S)$ stands for the set of nodes which have an ancestor in S (see Definition 2.1).

Hence, if γ is large, one of the terms of the product in Equation (38) must be large. Since the value of any sub-protocol is at most one, the numerator of each term cannot be large. It then must be the case that the denominator of at least one of those terms is close to zero, i.e., that u has a low-value ancestor controlled by A.⁴¹

The following key lemma formulates the above intuition, and shows that the biased-continuation attacker does not bias the original distribution of the protocol by too much, unless it has previously visited a low-value node controlled by A.

Lemma 4.6. *Let $\Pi = (A, B)$ be a protocol and let $A_{\Pi}^{(1)}$ be as in Algorithm 3.2. Then for every $\delta \in (0, \frac{1}{2}]$ there exists a constant $c = c(\delta) > 0$, such that for every $\delta' \geq \delta$ and $\gamma > 1$:*

$$\Pr_{\langle A_{\Pi}^{(1)}, B \rangle} \left[\text{desc} \left(\text{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right) \right] \leq \frac{2}{\gamma^c}.^{42}$$

Namely, the probability of reaching a γ -unbalanced node which does not have a δ' -low ancestor, for $\delta' \geq \delta$, is some inverse polynomial in γ . Looking ahead, we will apply this lemma for some $\gamma \in \text{poly}(n)$, where n is the security parameter given to the parties. At a high level, $\text{BiasedCont}^{\xi, \delta}$ gives a good (enough) approximation for BiasedCont when called on nodes that are at most $\text{poly}(n)$ -unbalanced. This lemma is useful since it gives a $1/\text{poly}(n)$ bound for the probability that $\text{BiasedCont}^{\xi, \delta}$ is called on nodes that are more than $\text{poly}(n)$ -unbalanced. Another important point is that the inverse polynomial (i.e., c) depends only on δ (and is independent of γ and δ'). This becomes crucial when analyzing the success probability of the approximated biased-continuation attacker.

Proof. The lemma is proven via the following three steps:

(1) Prove that for any such δ there exists $c > 0$, such that

$$\Pr_{\langle A^{(1)}, B \rangle} \left[\text{desc} \left(\text{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\text{Small}_{\Pi}^{\delta, A} \right) \right) \right] \leq \frac{2 - \text{val}(\Pi)}{\gamma^c} \quad (39)$$

for every $\gamma > 1$. Note that Equation (39) only considers descendants of $\text{Small}_{\Pi}^{\delta, A}$, and not proper descendants.

(2) Prove that for $\gamma > 1$ it holds that

$$\text{desc} \left(\text{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}} \left(\text{Small}_{\Pi}^{\delta, A} \right) \right) \subseteq \text{desc} \left(\text{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\text{Small}_{\Pi}^{\delta, A} \right) \right). \quad (40)$$

(3) Prove that for $\delta' > \delta$ it holds that

$$\text{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}} \left(\text{Small}_{\Pi}^{\delta', A} \right) \subseteq \text{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}} \left(\text{Small}_{\Pi}^{\delta, A} \right). \quad (41)$$

⁴¹This discussion is not entirely accurate, but it gives a good intuition for why unbalanced nodes relate to low-value ones. Indeed, the actual statement (Lemma 4.6) shows this discussion to hold only with high probability, which suffices for our needs.

⁴²Recall that for $S \subseteq \mathcal{V}(\Pi)$, $\overline{\text{desc}}(S)$ stands for the set of nodes which have an ancestor in S , but are not in S itself (see Definition 2.1).

It is clear that combining the above steps yields (a stronger version of) the lemma.

Proof of (1): Fix $\delta \in (0, \frac{1}{2}]$ and let $c := \alpha(\delta)$ be the value guaranteed in Lemma 2.18. The proof is by induction on the round complexity of Π .

Assume $\text{round}(\Pi) = 0$ and let ℓ be the single leaf of Π . By Definition 4.4, $\ell \notin \mathcal{UnBal}_\Pi^\gamma$ and thus $\mathcal{UnBal}_\Pi^\gamma = \emptyset$. Hence, for every $\delta > 0$,

$$\Pr_{\langle A^{(1)}, B \rangle} \left[\text{desc} \left(\mathcal{UnBal}_\Pi^\gamma \setminus \text{desc}(\text{Small}_\Pi^{\delta, A}) \right) \right] = \Pr_{\langle A^{(1)}, B \rangle} [\emptyset] = 0 \leq \frac{2 - \text{val}(\Pi)}{\gamma^c}.$$

Assume that Equation (39) holds for m -round protocols and that $\text{round}(\Pi) = m + 1$.

Assuming $e_{(A, B)}(\lambda, b) = 1$ for some $b \in \{0, 1\}$ (recall that λ denotes the empty string), then

$$\begin{aligned} \Pr_{\langle A^{(1)}, B \rangle} \left[\text{desc} \left(\mathcal{UnBal}_\Pi^\gamma \setminus \text{desc}(\text{Small}_\Pi^{\delta, A}) \right) \right] &= \Pr_{\langle (A^{(1)}, B)_b \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_b}^\gamma \setminus \text{desc}(\text{Small}_{\Pi_b}^{\delta, A}) \right) \right] \\ &= \Pr_{\langle A_{\Pi_b}^{(1)}, B_{\Pi_b} \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_b}^\gamma \setminus \text{desc}(\text{Small}_{\Pi_b}^{\delta, A}) \right) \right], \end{aligned}$$

where the second equality follows Proposition 3.5. The proof now follows from the induction hypothesis.

To complete the proof, we assume that $e_{(A, B)}(\lambda, b) \notin \{0, 1\}$ for both $b \in \{0, 1\}$, and let $p = e_{(A, B)}(\lambda, 0)$. The proof splits according to who controls the root of Π .

B controls $\text{root}(\Pi)$. We first prove that

$$\mathcal{UnBal}_\Pi^\gamma \setminus \text{desc} \left(\text{Small}_\Pi^{\delta, A} \right) = \left(\mathcal{UnBal}_{\Pi_0}^\gamma \setminus \text{desc} \left(\text{Small}_{\Pi_0}^{\delta, A} \right) \right) \cup \left(\mathcal{UnBal}_{\Pi_1}^\gamma \setminus \text{desc} \left(\text{Small}_{\Pi_1}^{\delta, A} \right) \right). \quad (42)$$

Let $u \in \mathcal{V}(\Pi)$. First, note that since B controls $\text{root}(\Pi)$, it holds that $e_{(A^{(1)}, B)}(\lambda, b) = e_{(A, B)}(\lambda, b)$, and thus, if $u \neq \text{root}(\Pi)$, it holds that $u \in \mathcal{UnBal}_\Pi^\gamma$ if and only if $u \in \mathcal{UnBal}_{\Pi_b}^\gamma$. Assume that $u \in \mathcal{UnBal}_\Pi^\gamma \setminus \text{desc} \left(\text{Small}_\Pi^{\delta, A} \right)$. Since $\gamma > 1$, it holds that $u \neq \text{root}(\Pi)$, and thus $u \in \mathcal{UnBal}_{\Pi_b}^\gamma$. Moreover, it follows that $u_1, \dots, u_{1, \dots, |u|} \notin \text{Small}_{\Pi_b}^{\delta, A}$, and thus $u \in \mathcal{UnBal}_{\Pi_b}^\gamma \setminus \text{desc} \left(\text{Small}_{\Pi_b}^{\delta, A} \right)$. For the other direction, assume $u \in \mathcal{UnBal}_{\Pi_b}^\gamma \setminus \text{desc} \left(\text{Small}_{\Pi_b}^{\delta, A} \right)$. As argued before, it holds that $u \in \mathcal{UnBal}_\Pi^\gamma$. Moreover, it follows that $u_1, \dots, u_{1, \dots, |u|} \notin \text{Small}_{\Pi_b}^{\delta, A}$, and since B controls $\text{root}(\Pi)$, it also holds that $\text{root}(\Pi) \notin \text{Small}_{\Pi_b}^{\delta, A}$. Hence, $u \in \mathcal{UnBal}_\Pi^\gamma \setminus \text{desc} \left(\text{Small}_\Pi^{\delta, A} \right)$. This complete the proof of Equation (42).

We can now write

$$\begin{aligned}
& \Pr_{\langle A^{(1)}, B \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi}^{\delta, A} \right) \right) \right] \\
&= e_{(A^{(1)}, B)}(\lambda, 0) \cdot \Pr_{\langle (A^{(1)}, B)_0 \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_0}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_0}^{\delta, A} \right) \right) \right] \\
&\quad + e_{(A^{(1)}, B)}(\lambda, 1) \cdot \Pr_{\langle (A^{(1)}, B)_1 \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_1}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_1}^{\delta, A} \right) \right) \right] \\
&= p \cdot \Pr_{\langle A_{\Pi_0}^{(1)}, B_{\Pi_0} \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_0}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_0}^{\delta, A} \right) \right) \right] \\
&\quad + (1-p) \cdot \Pr_{\langle A_{\Pi_1}^{(1)}, B_{\Pi_1} \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_1}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_1}^{\delta, A} \right) \right) \right] \\
&\leq p \cdot \frac{2 - \text{val}(\Pi_0)}{\gamma^c} + (1-p) \cdot \frac{2 - \text{val}(\Pi_1)}{\gamma^c} \\
&= \frac{2 - \text{val}(\Pi)}{\gamma^c}.
\end{aligned}$$

The first equality follows from Equation (42), the second equality follows from Proposition 3.5, and the inequality follows from the induction hypothesis.

A controls $\text{root}(\Pi)$. If $\text{val}(\Pi) \leq \delta$, then $\text{root}(\Pi) \in \mathcal{Small}_{\Pi}^{\delta, A}$. Therefore, $\mathcal{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi}^{\delta, A} \right) = \emptyset$ and the proof follows from a similar argument as in the base case.

In the complementary case, i.e., $\text{val}(\Pi) > \delta$, assume without loss of generality that $\text{val}(\Pi_0) \geq \text{val}(\Pi) \geq \text{val}(\Pi_1)$. We start with the case that $\text{val}(\Pi_1) > 0$ and the case that $\text{val}(\Pi_1) = 0$ is handled later. For $b \in \{0, 1\}$, let $\gamma_b := \frac{\text{val}(\Pi)}{\text{val}(\Pi_b)} \cdot \gamma$. By Claim 3.4, for $u \in \mathcal{V}(\Pi)$ with $u \neq \text{root}(\Pi)$ and $b = u_1$, it holds that

$$\frac{v_{(A^{(1)}, B)}(u)}{v_{(A, B)}(u)} = \frac{e_{(A, B)}(\lambda, b)}{e_{(A^{(1)}, B)}(\lambda, b)} \cdot \frac{v_{(A^{(1)}, B)_b}(u)}{v_{(A, B)_b}(u)} = \frac{\text{val}(\Pi_b)}{\text{val}(\Pi)} \cdot \frac{v_{(A^{(1)}, B)_b}(u)}{v_{(A, B)_b}(u)}.$$

Thus, $u \in \mathcal{UnBal}_{\Pi}^{\gamma}$ if and only if $u \in \mathcal{UnBal}_{\Pi_b}^{\gamma_b}$. Hence, using also the fact that $\text{root}(\Pi) \notin \mathcal{Small}_{\Pi}^{\delta, A}$ (since we assumed $\text{val}(\Pi) > \delta$), arguments similar to those used to prove Equation (42) yield that

$$\mathcal{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi}^{\delta, A} \right) = \left(\mathcal{UnBal}_{\Pi_0}^{\gamma_0} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_0}^{\delta, A} \right) \right) \cup \left(\mathcal{UnBal}_{\Pi_1}^{\gamma_1} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_1}^{\delta, A} \right) \right). \tag{43}$$

Moreover, we can write

$$\begin{aligned}
\Pr_{\langle (A^{(1)}, B)_b \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_b}^{\gamma_b} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_b}^{\delta, A} \right) \right) \right] &= \Pr_{\langle A_{\Pi_b}^{(1)}, B_{\Pi_b} \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_1}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_1}^{\delta, A} \right) \right) \right] \\
&\leq \frac{2 - \text{val}(\Pi_b)}{\gamma_b^c} \\
&= \left(\frac{\text{val}(\Pi_b)}{\text{val}(\Pi)} \right)^c \cdot \frac{2 - \text{val}(\Pi_b)}{\gamma^c}.
\end{aligned} \tag{44}$$

The first equality follows from Proposition 3.5, and the inequality follows from the induction hypothesis if $\gamma_b > 1$, and the fact that $\frac{2-\text{val}(\Pi_b)}{\gamma_b^c} \geq 1$ otherwise. We have that

$$\begin{aligned} & \Pr_{\langle A^{(1)}, B \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi}^{\delta, A} \right) \right) \right] \\ &= e_{\langle A^{(1)}, B \rangle}(\lambda, 0) \cdot \Pr_{\langle (A^{(1)}, B)_{0} \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_0}^{\gamma_0} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_0}^{\delta, A} \right) \right) \right] \\ &\quad + e_{\langle A^{(1)}, B \rangle}(\lambda, 1) \cdot \Pr_{\langle (A^{(1)}, B)_{1} \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_1}^{\gamma_1} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_1}^{\delta, A} \right) \right) \right] \\ &\leq p \cdot \left(\frac{\text{val}(\Pi_0)}{\text{val}(\Pi)} \right)^{1+c} \cdot \frac{2 - \text{val}(\Pi_0)}{\gamma^c} + (1-p) \cdot \left(\frac{\text{val}(\Pi_1)}{\text{val}(\Pi)} \right)^{1+c} \cdot \frac{2 - \text{val}(\Pi_1)}{\gamma^c}, \end{aligned} \quad (45)$$

where the equality follows from Equation (43), and the inequality follows from Equation (44) together with Claim 3.4. Letting $y = \frac{\text{val}(\Pi_0)}{\text{val}(\Pi)} - 1$, $x = \text{val}(\Pi)$ and $\lambda = \frac{p}{1-p}$, and noting that $\lambda y = \left(\frac{\text{val}(\Pi_0)}{\text{val}(\Pi)} - 1 \right) \cdot \frac{p}{1-p} = \frac{p \cdot \text{val}(\Pi_0) - p \cdot \text{val}(\Pi)}{\text{val}(\Pi) - p \cdot \text{val}(\Pi)} \leq \frac{p \cdot \text{val}(\Pi_0)}{\text{val}(\Pi)} \leq 1$, we can use Lemma 2.18 to deduce (after multiplying by $\frac{1-p}{\gamma^c}$) that

$$p \cdot \left(\frac{\text{val}(\Pi_0)}{\text{val}(\Pi)} \right)^{1+c} \cdot \frac{2 - \text{val}(\Pi_0)}{\gamma^c} + (1-p) \cdot \left(\frac{\text{val}(\Pi_1)}{\text{val}(\Pi)} \right)^{1+c} \cdot \frac{2 - \text{val}(\Pi_1)}{\gamma^c} \leq \frac{2 - \text{val}(\Pi)}{\gamma^c}, \quad (46)$$

completing the proof for the case $\text{val}(\Pi_1) > 0$.

It is left to argue the case that $\text{val}(\Pi_1) = 0$. In this case, according to Claim 3.4, it holds that $e_{\langle A^{(1)}, B \rangle}(\lambda, 0) = 1$ and $e_{\langle A^{(1)}, B \rangle}(\lambda, 1) = 0$. Hence, there are no unbalanced nodes in Π_1 , i.e., $\mathcal{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi}^{\delta, A} \right) \cap \mathcal{V}(\Pi_1) = \emptyset$. As before, let $\gamma_0 := \frac{\text{val}(\Pi)}{\text{val}(\Pi_0)} \cdot \gamma = p \cdot \gamma$ (The latter equality holds since $\text{val}(\Pi) = p \cdot \text{val}(\Pi_0)$.) Arguments similar to those used to prove Equation (43) yield that

$$\mathcal{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi}^{\delta, A} \right) = \mathcal{UnBal}_{\Pi_0}^{\gamma_0} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_0}^{\delta, A} \right). \quad (47)$$

It follows that

$$\begin{aligned} & \Pr_{\langle A^{(1)}, B \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi}^{\delta, A} \right) \right) \right] \\ &= e_{\langle A^{(1)}, B \rangle}(\lambda, 0) \cdot \Pr_{\langle (A^{(1)}, B)_{0} \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi_0}^{\gamma_0} \setminus \text{desc} \left(\mathcal{Small}_{\Pi_0}^{\delta, A} \right) \right) \right] \\ &\leq \left(\frac{1}{p} \right)^{1+c} \cdot \frac{2 - \text{val}(\Pi_0)}{\gamma^c}. \end{aligned}$$

Applying Lemma 2.18 with the same parameters as above completes the proof.

Proof of (2): Fix $\gamma > 1$. We prove that

$$\text{frnt} \left(\mathcal{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}} \left(\mathcal{Small}_{\Pi}^{\delta, A} \right) \right) \subseteq \mathcal{UnBal}_{\Pi}^{\gamma} \setminus \text{desc} \left(\mathcal{Small}_{\Pi}^{\delta, A} \right), \quad (48)$$

which clearly derives (2).

⁴³Recall that for a set $\mathcal{S} \subset \mathcal{V}(\Pi)$, $\text{frnt}(\mathcal{S})$ stands for the frontier of \mathcal{S} , i.e., the set of nodes belong to \mathcal{S} , whose ancestors do not belong to \mathcal{S} (see Definition 2.1).

Let $u \in \text{frnt} \left(\text{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}} \left(\text{Small}_{\Pi}^{\delta, A} \right) \right)$. We prove Equation (48) by showing that $u \notin \text{Small}_{\Pi}^{\delta, A}$. Since $\gamma > 1$ and $u \in \text{UnBal}_{\Pi}^{\gamma}$, it is clear that $u \neq \text{root}(\Pi)$. Let w be the parent of u . By the choice of u , it follows that $w \notin \text{UnBal}_{\Pi}^{\gamma}$, and thus $v_{(A^{(1)}, B)}(w) < \gamma \cdot v_{(A, B)}(w)$. We write

$$\begin{aligned} \gamma \cdot v_{(A, B)}(w) \cdot e_{(A^{(1)}, B)}(w, u) &> v_{(A^{(1)}, B)}(w) \cdot e_{(A^{(1)}, B)}(w, u) \\ &= v_{(A^{(1)}, B)}(u) \\ &\geq \gamma \cdot v_{(A, B)}(u) \\ &= \gamma \cdot v_{(A, B)}(w) \cdot e_{(A, B)}(w, u). \end{aligned} \tag{49}$$

We conclude that $e_{(A, B)}(w, u) < e_{(A^{(1)}, B)}(w, u)$, and thus it must be the case that w is controlled by A. By Claim 3.4, it holds that $e_{(A^{(1)}, B)}(w, u) = e_{(A, B)}(w, u) \cdot \frac{\text{val}(\Pi_u)}{\text{val}(\Pi_w)}$, and thus $\text{val}(\Pi_u) > \text{val}(\Pi_w)$. Finally, observe that $w \notin \text{Small}_{\Pi}^{\delta, A}$, since otherwise $u \in \overline{\text{desc}} \left(\text{Small}_{\Pi}^{\delta, A} \right)$. It follows that $\text{val}(\Pi_w) > \delta$, and hence $\text{val}(\Pi_u) > \delta$, as required.

Proof of (3): Note that for every $\delta' \geq \delta$ it holds that $\text{Small}_{\Pi}^{\delta, A} \subseteq \text{Small}_{\Pi}^{\delta', A}$. Hence, $\text{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}}(\text{Small}_{\Pi}^{\delta', A}) \subseteq \text{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}}(\text{Small}_{\Pi}^{\delta, A})$, and the proof follows. \square

Lemma 4.6 allows us to bound the probability that the (ideal) biased-continuation attacker hits unbalanced nodes with the probability that the *original* protocol hits A-controlled low-value nodes. Indeed, consider the first time $(A^{(1)}, B)$ reaches a γ -unbalanced node u . If this process generates an A-controlled low-value ancestor for u , then this ancestor cannot be γ -unbalanced, and thus the probability of hitting it (and in turn hitting u) is bounded by γ times the probability of the original protocol hitting A-controlled low-value nodes. In the complementary case, in which no A-controlled low-value node was generated before hitting u , then the probability of hitting u is bounded by Lemma 4.6. The above discussion is stated formally next.

Corollary 4.7. *Let $\Pi = (A, B)$ be a protocol, let $\delta \in (0, \frac{1}{2}]$, and let $c = c(\delta)$ be according Lemma 4.6. Then*

$$\Pr_{\langle A^{(1)}, B \rangle} \left[\text{desc} \left(\text{UnBal}_{\Pi}^{\gamma} \right) \right] \leq \gamma \cdot \Pr_{\langle A, B \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right] + \frac{2}{\gamma^c},$$

for any $\delta' \geq \delta$ and $\gamma > 1$.

Proof. Our first step is to show that

$$\text{desc} \left(\text{UnBal}_{\Pi}^{\gamma} \right) \subseteq \text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \setminus \text{UnBal}_{\Pi}^{\gamma} \right) \cup \text{desc} \left(\text{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right). \tag{50}$$

Indeed, let $\ell \in \text{desc} \left(\text{UnBal}_{\Pi}^{\gamma} \right)$, and let $u \in \text{frnt} \left(\text{UnBal}_{\Pi}^{\gamma} \right)$ such that $\ell \in \text{desc}(u)$. Assume that u has an A-controlled δ' -value ancestor w , i.e., that there exists $w \in \text{Small}_{\Pi}^{\delta', A}$ such that $u \in \overline{\text{desc}}(w)$. Then by the choice of u , w must be γ -balanced, i.e., $w \notin \text{UnBal}_{\Pi}^{\gamma}$. It follows that $\ell \in \text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \setminus \text{UnBal}_{\Pi}^{\gamma} \right)$. In the complementary case, i.e., that $u \notin \overline{\text{desc}} \left(\text{Small}_{\Pi}^{\delta', A} \right)$, it holds that $\ell \in \text{desc} \left(\text{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc}} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right)$. This yields Equation (50).

We can now compute

$$\begin{aligned} \Pr_{\langle A^{(1)}, B \rangle} [\text{desc}(\mathcal{UnBal}_{\Pi}^{\gamma})] &\leq \Pr_{\langle A^{(1)}, B \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \setminus \mathcal{UnBal}_{\Pi}^{\gamma} \right) \right] \\ &\quad + \Pr_{\langle A^{(1)}, B \rangle} \left[\text{desc} \left(\mathcal{UnBal}_{\Pi}^{\gamma} \setminus \overline{\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \right)} \right) \right] \\ &\leq \gamma \cdot \Pr_{\langle A, B \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right] + \frac{2}{\gamma^c}, \end{aligned}$$

where the second inequality follows from the definition of $\mathcal{UnBal}_{\Pi}^{\gamma}$ and Lemma 4.6. \square

Finally, using Corollary 4.7, we can derive the main conclusion of Section 4.1 — the approximated biased-continuation attacker successfully biases protocols in which the probability of hitting A-controlled low-value nodes is small.

Corollary 4.8. *Let $\Pi = (A, B)$ be an m -round protocol, let $\delta \in (0, \frac{1}{2}]$, and let $c = c(\delta)$ be according to Lemma 4.6. Then*

$$\text{SD} \left(\langle A_{\Pi}^{(1)}, B \rangle, \langle A_{\Pi}^{(1, \xi, \delta')}, B \rangle \right) \leq 2 \cdot m \cdot \gamma \cdot \left(\xi + \Pr_{\langle A, B \rangle} \left[\text{desc}(\text{Small}_{\Pi}^{\delta', A}) \right] \right) + \frac{2}{\gamma^c}$$

for any $\delta' \geq \delta$, $\xi > 0$ and $\gamma > 1$.

Proof. The proof immediately follows by plugging Corollary 4.7 into Lemma 4.5. \square

Bounding the Probability of Hitting Low-Density Sets. Our final statement in Section 4.1 is a generalization of Corollary 4.7 to arbitrary sets of nodes (i.e., not only unbalanced) and to the recursive version of the ideal biased-continuation attacker. This generalization will be helpful in the rest of the section.

Proposition 4.9. *Let $\Pi = (A, B)$ be a protocol, let $\delta \in (0, \frac{1}{2}]$, and let $c = c(\delta)$ be according to Lemma 4.6. Then the following holds for any $\delta' \geq \delta$:*

1. For every $k \in \mathbb{N}$ and any $\gamma_1, \dots, \gamma_k > 1$ it holds that

$$\Pr_{\langle A_{\Pi}^{(k)}, B \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right] \leq \Pr_{\langle A, B \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right] \cdot \prod_{i=1}^k \gamma_i + 2 \cdot \sum_{i=1}^k \frac{\prod_{j=i+1}^k \gamma_j}{\gamma_i^c}.$$

2. For every $\mathcal{S} \subseteq \mathcal{V}(\Pi)$ with $\Pr_{\langle A, B \rangle} [\text{desc}(\mathcal{S})] \leq \alpha$, any $\beta \geq \Pr_{\langle A, B \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right]$, every $k \in \mathbb{N}$ and any $\gamma = (\gamma_1, \dots, \gamma_k)$ with $\gamma_i > 1$ for all $i \in [k]$, it holds that

$$\Pr_{\langle A_{\Pi}^{(k)}, B \rangle} [\text{desc}(\mathcal{S})] \leq \phi^{\text{Bal}}(\alpha, \beta, \delta, \gamma) := (\alpha + \beta) \cdot \prod_{i=1}^k \gamma_i + 2 \cdot \sum_{i=1}^k \frac{\prod_{j=i+1}^k \gamma_j}{\gamma_i^c}.$$

Proof. The proof follows from an analysis similar to the proof of Corollary 4.7. \square

4.2 Attacking Pruned Protocols

In the previous section we showed that if the probability to visit A-controlled low-value nodes is small, the approximated biased-continuation attacker biases the protocol almost as well as its ideal variant does. For some protocols, however, this probability might be arbitrarily large, so the analysis in Section 4.1 does not suffice to argue that the approximated biased-continuation attacker successfully biases *any* protocol. For an arbitrary protocol, however, we can define a pruned variant of it, such that the probability of hitting A-controlled low-value nodes is indeed small. The above corollary yields that the approximated biased-continuation attacker successfully biases the above variant.

The definition of these pruned variants and the analysis of attacking them using the approximated biased-continuation attacker are the focus of this section. We show that the (recursive variant) of the approximated biased-continuation attacker defined in the previous section is very effective against the (approximated) *pruned* variant of the protocol, defined next.

4.2.1 Pruned Protocols

In the pruned variant of a protocol $\Pi = (A, B)$, the edge distribution remains intact, while the controlling scheme is changed, giving the control to B on low-value nodes, and to A on high-value nodes.

Definition 4.10 (the pruned variant of a protocol). *Let $\Pi = (A, B)$ be an m -round protocol and let $\delta \in (0, \frac{1}{2})$. In the δ -pruned variant of Π , denoted by $\Pi^{[\delta]} = (A_{\Pi}^{[\delta]}, B_{\Pi}^{[\delta]})$, the parties follow the protocol Π , where $A_{\Pi}^{[\delta]}$ and $B_{\Pi}^{[\delta]}$ take the roles of A and B respectively, with the following exception occurring the first time the protocol's transcript u is in $\text{Small}_{\Pi}^{\delta} \cup \text{Large}_{\Pi}^{\delta}$:*

If $u \in \text{Large}_{\Pi}^{\delta}$, set $C = A_{\Pi}^{[\delta]}$; otherwise set $C = B_{\Pi}^{[\delta]}$. The party C takes control of the node u , samples a leaf $\ell \leftarrow \langle \Pi_u \rangle$, and then, bit by bit, sends $\ell_{|u|+1, \dots, m}$ to the other party.

Namely, the first time the value of the protocol is close to either 1 or 0, the party interested in this value (i.e., $A_{\Pi}^{[\delta]}$ for 1, and $B_{\Pi}^{[\delta]}$ for 0) takes control and decides the outcome (without changing the value of the protocol). Hence, the protocol is effectively pruned at these leaves (each such a node is effectively a parent of two leaves).

For every protocol Π , its pruned variant $\Pi^{[\delta]}$ is a well-defined protocol, so the analysis of Section 3 can be applied.⁴⁴ (Later, our almost-final attacker will “pretend” it is actually running on this pruned variant, rather than on the original protocol). The pruned variant of a protocol, however, might *not* be efficiently computed, even if OWFs do not exist. To cope with this efficiency issue, we consider an approximated variant of the pruned protocol.

4.2.1.1 Approximated Pruned Protocols

To define the approximated pruned protocols, we begin by defining two algorithms, both of which can be efficiently implemented assuming OWFs do not exist for an appropriate set of parameters. The first algorithm samples an honest (i.e., unbiased) continuation of the protocol.

⁴⁴Note that in the pruned protocol, the parties' turns might not alternate (i.e., the same party might send several consecutive bits), even if they do alternate in the original protocol. Rather, the protocol's control scheme (determining which party is active at a given point) is a function of the protocol's transcript and the original protocol's control scheme. Such schemes are consistent with the ones considered in the previous sections.

Definition 4.11 (approximated honest continuation). Let Π be an m -round protocol, and let HonCont_Π be the algorithm that on node $u \in \mathcal{V}(\Pi)$ returns $\ell \leftarrow \langle \Pi_u \rangle$. Algorithm HC is a ξ -Honest-Continuator for Π , if $\Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1): \text{SD}(\text{HC}(\ell_{1,\dots,i}), \text{HonCont}_\Pi(\ell_{1,\dots,i})) > \xi] \leq \xi$. Let HonCont_Π^ξ be an arbitrary (but fixed) ξ -honest-continuator for Π .

The second algorithm estimates the value of a given transcript (i.e., a node) of the protocol.

Definition 4.12 (estimator). Let Π be an m -round protocol. A deterministic algorithm Est is a ξ -Estimator for Π , if $\Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1): |\text{Est}(\ell_{1,\dots,i}) - \text{val}(\Pi_{\ell_{1,\dots,i}})| > \xi] \leq \xi$. Let Est_Π^ξ be an arbitrary (but fixed) ξ -estimator for Π .

Using the above estimator, we define the approximated version of the low and high value nodes.

Definition 4.13 (approximated low-value and high-value nodes). For protocol Π , $\delta \in (0, \frac{1}{2})$ and a deterministic real-value algorithm Est , let

- $\text{Small}_\Pi^{\delta, \text{Est}} = \{u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi) : \text{Est}(u) \leq \delta\}$;
- $\text{Large}_\Pi^{\delta, \text{Est}} = \{u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi) : \text{Est}(u) \geq 1 - \delta\}$.

For $\xi \in [0, 1]$, let $\text{Small}_\Pi^{\delta, \xi} = \text{Small}_\Pi^{\delta, \text{Est}_\Pi^\xi}$.

We can now define the approximated pruned protocol, which is the oracle variant of the ideal pruned protocol.

Definition 4.14 (the approximated pruned variant of a protocol). Let $\Pi = (A, B)$ be an m -round protocol, let $\delta \in (0, \frac{1}{2})$, let HC be an algorithm, and let Est be a deterministic real value algorithm. The $(\delta, \text{Est}, \text{HC})$ -approximately pruned variant of Π , denoted $\Pi^{[\delta, \text{Est}, \text{HC}]} = \left(A_\Pi^{[\delta, \text{Est}, \text{HC}]}, B_\Pi^{[\delta, \text{Est}, \text{HC}]} \right)$, is defined as follows.

Control Scheme: the parties follow the control scheme of the protocol Π , where $A_\Pi^{[\delta, \text{Est}, \text{HC}]}$ and $B_\Pi^{[\delta, \text{Est}, \text{HC}]}$ take the roles of A and B respectively, with the following exception occurring the first time the protocol's transcript u is in $\text{Small}_\Pi^{\delta, \text{Est}} \cup \text{Large}_\Pi^{\delta, \text{Est}}$: if $u \in \text{Large}_\Pi^{\delta, \text{Est}}$ set $C = A_\Pi^{[\delta, \text{Est}, \text{HC}]}$; otherwise set $C = B_\Pi^{[\delta, \text{Est}, \text{HC}]}$. The party C takes control of all nodes in $\text{desc}(u)$ (i.e., nodes for which u is an ancestor).

Execution: for a protocol's transcript u and a party C who controls u , C sets $\ell = \text{HC}(u)$ and sends $\ell_{|u|+1}$ to the other party.⁴⁵

For $\delta \in (0, \frac{1}{2})$ and $\xi, \xi' \in [0, 1]$, let $\Pi^{[\delta, \xi, \xi']} = \Pi^{[\delta, \text{Est}_\Pi^\xi, \text{HonCont}_\Pi^{\xi'}]}$ and $\Pi^{[\delta, \xi]} = \Pi^{[\delta, \xi, \xi]}$, and the same notation is used for the parties of the pruned protocol.

Namely, in $\Pi^{[\delta, \xi]}$, the parties follow the control scheme of Π until reaching a node in $\text{Small}_\Pi^{\delta, \xi} \cup \text{Large}_\Pi^{\delta, \xi}$ for the first time. Upon reaching such a node, the control moves to (and stays with) A if $u \in \text{Large}_\Pi^{\delta, \xi}$, or B if $u \in \text{Small}_\Pi^{\delta, \xi}$. The fact that the messages sent by the parties are determined by the answers of HonCont_Π^ξ , instead of by their random coins, makes them *stateless* throughout the execution of the protocol. This will be crucial when implementing our final attacker.

⁴⁵This happens to every transcript, even those that are not children of $\text{Small}_\Pi^{\delta, \text{Est}} \cup \text{Large}_\Pi^{\delta, \text{Est}}$.

Properties of the approximated pruned protocol. We now prove some important properties of the approximated pruned protocol. The first property is that it is a good approximation of the ideal pruned protocol.

Lemma 4.15. *Let $\Pi = (A, B)$ be an m -round protocol. Then*

$$\text{SD} \left(\langle \Pi \rangle, \langle \Pi^{[\delta, \xi]} \rangle \right) \leq 2 \cdot m \cdot \xi$$

for every $\delta \in (0, \frac{1}{2}]$ and $\xi \in (0, 1)$.

Note that the leaf distributions of Π and $\Pi^{[\delta]}$ are identical, so the above lemma indeed shows that the leaf distributions of the ideal and approximated pruned protocols are close. Moreover, the above bound is a simple implication of the approximation guarantee of the honest-continuator, and does not depend on δ .

Proof. By definition, every message in $\Pi^{[\delta, \xi, 0]}$ is set by calling a perfect honest-continuator for Π . Thus, $\langle \Pi \rangle \equiv \langle \Pi^{[\delta, \xi, 0]} \rangle$, and it suffices to bound $\text{SD} \left(\langle \Pi^{[\delta, \xi, 0]} \rangle, \langle \Pi^{[\delta, \xi]} \rangle = \langle \Pi^{[\delta, \xi, \xi]} \rangle \right)$, which we do by applying Lemma 2.14.

For function O , let H^O be an algorithm that outputs the transcript of a random execution of $\Pi^{[\delta, \text{Est}_{\Pi}^{\xi}, O]}$. Let f and g be the (random) functions HonCont_{Π} and $\text{HonCont}_{\Pi}^{\xi}$ respectively, with the exception that $f(\perp) = g(\perp) = \perp$. By construction, it holds that

$$\text{SD} \left(\langle \Pi^{[\delta, \xi, 0]} \rangle, \langle \Pi^{[\delta, \xi, \xi]} \rangle \right) = \text{SD} \left(H^f, H^g \right).$$

For $i \in [m]$, let D_i be i 'th node in a random execution of Π (such a node consists of $i - 1$ bits), and let $\mathcal{FailCont}_{\Pi}^{\xi, i} = \left\{ u \in \mathcal{V}(\Pi) : |u| = i - 1 \wedge \text{SD} \left(\text{HonCont}_{\Pi}(u), \text{HonCont}_{\Pi}^{\xi}(u) \right) > \xi \right\}$. By definition,

$$\begin{aligned} \Pr_{u \leftarrow D_i} \left[u \in \mathcal{FailCont}_{\Pi}^{\xi, i} \right] &= \Pr_{\ell \leftarrow (\Pi)} \left[\text{SD} \left(\text{HonCont}(\ell_{1, \dots, i-1}), \text{HonCont}_{\Pi}^{\xi}(\ell_{1, \dots, i-1}) \right) > \xi \right] \\ &\leq \Pr_{\ell \leftarrow (\Pi)} \left[\exists i \in [m] : \text{SD} \left(\text{HonCont}(\ell_{1, \dots, i-1}), \text{HonCont}_{\Pi}^{\xi}(\ell_{1, \dots, i-1}) \right) > \xi \right] \\ &\leq \xi, \end{aligned}$$

and thus

$$\begin{aligned} &E_{u \leftarrow D_i} [\text{SD}(f(u), g(u))] \\ &= E_{u \leftarrow D_i} \left[\text{SD} \left(\text{HonCont}_{\Pi}(u), \text{HonCont}_{\Pi}^{\xi}(u) \right) \right] \\ &= \Pr_{u \leftarrow D_i} \left[u \in \mathcal{FailCont}_{\Pi}^{\xi, i} \right] \cdot E_{u \leftarrow D_i} \left[\text{SD} \left(\text{HonCont}_{\Pi}(u), \text{HonCont}_{\Pi}^{\xi}(u) \right) \mid u \in \mathcal{FailCont}_{\Pi}^{\xi, i} \right] \\ &\quad + \Pr_{u \leftarrow D_i} \left[u \notin \mathcal{FailCont}_{\Pi}^{\xi, i} \right] \cdot E_{u \leftarrow D_i} \left[\text{SD} \left(\text{HonCont}_{\Pi}(u), \text{HonCont}_{\Pi}^{\xi}(u) \right) \mid u \notin \mathcal{FailCont}_{\Pi}^{\xi, i} \right] \\ &\leq \xi + \xi = 2\xi, \end{aligned}$$

where the first equality follows since $D_i(\perp) = 0$.

Let Q_i denote the i 'th query to f in a random execution of H^f (note that by construction, such a query always exists) and let $Q = (Q_1, \dots, Q_m)$. By definition, $Q_i \equiv D_i$, and thus

$$\Pr_{(q_1, \dots, q_m) \leftarrow Q} [\exists i \in [m] : q_i \neq \perp \wedge Q_i(q_i) > D_i(q_i)] = 0.$$

The proof follows by Lemma 2.14, letting $k = m$, $a = 2\xi$, $\lambda = 1$ and $b = 0$. \square

The second (and most important) property of the approximated pruned protocol is that it visits low-value nodes under A 's control only with small probability. While it is impossible to reach such a node in the *ideal* pruned protocol, bounding this probability is not an immediate corollary of Lemma 4.15. This is because the value of each node in both protocols might not be the same, and because the control scheme of these protocols might be different. It turns out that the bound for the above probability depends on the probability of the original protocol visiting nodes whose value is close to the pruning value.

Definition 4.16. For protocol Π , $\xi \in (0, 1)$ and $\delta \in (0, \frac{1}{2})$, let

$$\mathcal{B}order_{\Pi}^{\delta, \xi} = \{u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi) : \text{val}(\Pi_u) \in (\delta \pm \xi] \vee \text{val}(\Pi_u) \in [1 - \delta \pm \xi)\},$$

and let $\text{border}_{\Pi}(\delta, \xi) = \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\mathcal{B}order_{\Pi}^{\delta, \xi} \right) \right]$.

Namely, $\mathcal{B}order_{\Pi}^{\delta, \xi}$ are those nodes that are ξ -close to the “border” between $\mathcal{S}mall_{\Pi}^{\delta} \cup \mathcal{L}arge_{\Pi}^{\delta}$ and the rest of the nodes.

Lemma 4.17. Let $\delta \in (0, \frac{1}{2}]$, let $\varepsilon \in (0, \delta)$, let $\xi \in (0, 1)$ and let $\tilde{\Pi} = (\tilde{A}, \tilde{B}) = \Pi^{[\delta, \xi]}$ be the (δ, ξ) -approximately pruned variant of an m -round protocol Π . Then

$$\Pr_{\langle \tilde{\Pi} \rangle} \left[\text{desc} \left(\mathcal{S}mall_{\tilde{\Pi}}^{\delta - \varepsilon, \tilde{A}} \right) \right] \leq \text{border}_{\Pi}(\delta, \xi) + \frac{4 \cdot m \cdot \xi}{\varepsilon}.$$

Namely, as long as the probability of reaching nodes whose value is ξ -close to δ is small in the *original* protocol, the probability of hitting low-value nodes in the approximated pruned protocol is small as well. For proving Lemma 4.17, we use the following proposition, showing that if two protocols are close and there exists a set of nodes whose value (the probability that the common output is one conditioned on reaching these nodes) is large in one protocol but small in the other, then the probability of reaching this set is small.

Proposition 4.18. Let $\Pi = (A, B)$ and $\Pi' = (C, D)$ be two m -round protocols with $\chi_{\Pi} \equiv \chi_{\Pi'}$, and let $\mathcal{F} \subseteq \mathcal{V}(\Pi)$ be a frontier. Assume that $\text{SD}(\langle \Pi \rangle, \langle \Pi' \rangle) \leq \varepsilon$, that $\Pr_{\langle \Pi \rangle} [\mathcal{L}_1(\Pi) \mid \text{desc}(\mathcal{F})] \leq \alpha$, and that $\Pr_{\langle \Pi' \rangle} [\mathcal{L}_1(\Pi) \mid \text{desc}(\mathcal{F})] \geq \beta$, for some $\varepsilon > 0$ and $0 \leq \alpha < \beta \leq 1$. Then, $\Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F})] \leq \varepsilon \cdot \frac{1 + \beta}{\beta - \alpha}$.

Note that since both Π and Π' have m -rounds, it holds that $\mathcal{V}(\Pi) = \mathcal{V}(\Pi')$ and $\mathcal{L}(\Pi) = \mathcal{L}(\Pi')$. Moreover, since $\chi_{\Pi} \equiv \chi_{\Pi'}$, it also holds that $\mathcal{L}_1(\Pi)$, the set of 1-leaves in Π , is identical to $\mathcal{L}_1(\Pi')$, the set of 1-leaves in Π' .

Proof. Let $\mu = \Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F})]$, $\mu' = \Pr_{\langle \Pi' \rangle} [\text{desc}(\mathcal{F})]$ and $\mathcal{S} = \{\ell \in \mathcal{L}_1(\Pi) : \chi_{\Pi}(\ell) = 1\}$. It follows that

$$\Pr_{\langle \Pi \rangle} [\mathcal{S}] = \Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F})] \cdot \Pr_{\langle \Pi \rangle} [\mathcal{L}_1(\Pi) \mid \text{desc}(\mathcal{F})] \leq \mu \cdot \alpha \tag{51}$$

and that

$$\Pr_{\langle \Pi' \rangle} [\mathcal{S}] = \Pr_{\langle \Pi' \rangle} [\text{desc}(\mathcal{F})] \cdot \Pr_{\langle \Pi' \rangle} [\mathcal{L}_1(\Pi) \mid \text{desc}(\mathcal{F})] \geq \mu' \cdot \beta. \tag{52}$$

Moreover, since $\text{SD}(\langle \Pi \rangle, \langle \Pi' \rangle) \leq \varepsilon$, it follows that $\mu' \geq \mu - \varepsilon$ and that $\Pr_{\langle \Pi' \rangle}[\mathcal{S}] - \Pr_{\langle \Pi \rangle}[\mathcal{S}] \leq \varepsilon$. Putting it all together, we get

$$\begin{aligned} \varepsilon &\geq \Pr_{\langle \Pi' \rangle}[\mathcal{S}] - \Pr_{\langle \Pi \rangle}[\mathcal{S}] \\ &\geq \mu' \cdot \beta - \mu \cdot \alpha \\ &\geq (\mu - \varepsilon) \cdot \beta - \mu \cdot \alpha \\ &= (\beta - \alpha) \cdot \mu - \beta \cdot \varepsilon, \end{aligned}$$

which implies the proposition. \square

Proof of Lemma 4.17. Let $\mathcal{F}\text{ailEst}_{\Pi}^{\xi} = \left\{ u \in \mathcal{V}(\Pi) : \left| \text{val}(\Pi_u) - \text{Est}_{\Pi}^{\xi}(u) \right| > \xi \right\}$ and let $\mathcal{F} = \text{frnt} \left(\text{Small}_{\tilde{\Pi}}^{\delta - \varepsilon, \tilde{A}} \right) \setminus \left(\text{Border}_{\tilde{\Pi}}^{\delta, \xi} \cup \mathcal{F}\text{ailEst}_{\tilde{\Pi}}^{\xi} \right)$. It follows that

$$\Pr_{\langle \tilde{\Pi} \rangle} \left[\text{desc} \left(\text{Small}_{\tilde{\Pi}}^{\delta - \varepsilon, \tilde{A}} \right) \right] \leq \Pr_{\langle \tilde{\Pi} \rangle} \left[\text{desc} \left(\text{Border}_{\tilde{\Pi}}^{\delta, \xi} \cup \mathcal{F}\text{ailEst}_{\tilde{\Pi}}^{\xi} \right) \right] + \Pr_{\langle \tilde{\Pi} \rangle} [\text{desc}(\mathcal{F})]. \quad (53)$$

By Lemma 4.15, it holds that

$$\Pr_{\langle \tilde{\Pi} \rangle} \left[\text{desc} \left(\text{Border}_{\tilde{\Pi}}^{\delta, \xi} \cup \mathcal{F}\text{ailEst}_{\tilde{\Pi}}^{\xi} \right) \right] \leq \text{border}_{\tilde{\Pi}}(\delta, \xi) + 3 \cdot m \cdot \xi. \quad (54)$$

Let $u \in \mathcal{F}$. Since u is under \tilde{A} 's control, we have $\text{Est}_{\tilde{\Pi}}^{\xi}(u) > \delta$. Since $u \notin \mathcal{F}\text{ailEst}_{\tilde{\Pi}}^{\xi}$, we have $\text{val}(\Pi_u) > \delta - \xi$, and since $u \notin \text{Border}_{\tilde{\Pi}}^{\delta, \xi}$, we have $\text{val}(\Pi_u) \geq \delta + \xi$. By definition, $\text{val}(\tilde{\Pi}_u) \leq \delta - \varepsilon$. Thus, $\Pr_{\langle \tilde{\Pi} \rangle} [\mathcal{L}_1(\Pi) \mid \text{desc}(\mathcal{F})] \leq \delta - \varepsilon$ and $\Pr_{\langle \Pi \rangle} [\mathcal{L}_1(\Pi) \mid \text{desc}(\mathcal{F})] \geq \delta + \xi$. Finally, by Proposition 4.18 and Lemma 4.15,

$$\Pr_{\langle \tilde{\Pi} \rangle} [\text{desc}(\mathcal{F})] \leq 2 \cdot m \cdot \xi \cdot \frac{1 + \delta - \varepsilon}{\xi + \varepsilon} \leq \frac{4 \cdot m \cdot \xi}{\varepsilon}. \quad (55)$$

Plugging Equations (54) and (55) into Equation (53) completes the proof of the lemma. \square

Lemma 4.17 upper-bounds the probability of visiting a low-value node under A 's control in the approximated pruned protocol with the probability of visiting nodes whose value is close to the point of pruning in the original protocol. Given a protocol and a pruning point, the latter probability might be large. We argue, however, that if we allow a small deviation from the point of pruning, this probability is small.

Proposition 4.19. *Let Π be an m -round protocol, let $\delta \in (0, \frac{1}{2}]$, and let $\xi \in (0, 1)$. If $\xi \leq \frac{\delta^2}{16m^2}$, then there exists $j \in \mathcal{J} := \{0, 1, \dots, \lceil m/\sqrt{\xi} \rceil\}$ such that $\text{border}_{\Pi}(\delta', \xi) \leq m \cdot \sqrt{\xi}$ for $\delta' = \delta/2 + j \cdot 2\xi \in [\frac{\delta}{2}, \delta]$.*

Proof. For $j \in \mathcal{J}$, let $\delta'(j) = \delta/2 + j \cdot 2\xi$. From the definition of \mathcal{J} , it is clear that $\delta'(j) \in [\frac{\delta}{2}, \delta]$ for every $j \in \mathcal{J}$. Hence, it is left to argue that $\exists j \in \mathcal{J}$ such that $\text{border}_{\Pi}(\delta'(j), \xi) \leq m \cdot \sqrt{\xi}$.

For $i \in [m]$, let $\text{Border}_{\Pi}^{\delta, \xi, i} = \left\{ u \in \mathcal{V}(\Pi) : u \in \text{Border}_{\Pi}^{\delta, \xi} \wedge |u| = i - 1 \right\}$. It holds that

$$\begin{aligned} \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Border}_{\Pi}^{\delta, \xi} \right) \right] &\leq \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\bigcup_{i \in [m]} \text{Border}_{\Pi}^{\delta, \xi, i} \right) \right] \\ &\leq \sum_{i=1}^m \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Border}_{\Pi}^{\delta, \xi, i} \right) \right]. \end{aligned} \quad (56)$$

For every $i \in [m]$, let $\mathcal{N}(i) = \{j \in \mathcal{J} : \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Border}_{\Pi}^{\delta'(j), \xi, i} \right) > \sqrt{\xi} \right]\}$ and let $\mathcal{N} = \cup_{i \in [m]} \mathcal{N}(i)$. We use the following claim.

Claim 4.20. *It holds that $|\mathcal{N}(i)| < 1/\sqrt{\xi}$ for every $i \in [m]$.*

Proof of Claim 4.20. Assume towards a contradiction that there exists $i \in [m]$ such that $|\mathcal{N}(i)| \geq 1/\sqrt{\xi}$. Let D_i be the distribution over $\{0, 1\}^i$, described by outputting ℓ_i , for $\ell \leftarrow \langle \Pi \rangle$. We get that $\Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Border}_{\Pi}^{\delta'(j), \xi, i} \right) \right] = D_i \left(\text{Border}_{\Pi}^{\delta'(j), \xi, i} \right)$. Since $\text{Border}_{\Pi}^{\delta'(j), \xi, i} \cap \text{Border}_{\Pi}^{\delta'(j'), \xi, i} = \emptyset$ for every $j \neq j' \in \mathcal{J}$, it holds that

$$\begin{aligned} 1 &\geq \sum_{j \in \mathcal{J}} D_i \left(\text{Border}_{\Pi}^{\delta'(j), \xi, i} \right) \\ &\geq \sum_{j \in \mathcal{N}(i)} D_i \left(\text{Border}_{\Pi}^{\delta'(j), \xi, i} \right) \\ &> |\mathcal{N}(i)| \cdot \sqrt{\xi} \geq 1, \end{aligned}$$

and a contradiction is derived, where the last inequality follows the assumption that $|\mathcal{N}(i)| \geq 1/\sqrt{\xi}$. \square

Claim 4.20 yields that $|\mathcal{N}| \leq \sum_{i=1}^m |\mathcal{N}(i)| < \frac{m}{\sqrt{\xi}} < |\mathcal{J}|$. Thus, $\exists j \in \mathcal{J}$ such that $j \notin \mathcal{N}$. Set $\delta' = \delta'(j)$. It holds that $\Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Border}_{\Pi}^{\delta', \xi, i} \right) \right] \leq \sqrt{\xi}$ for every $i \in [m]$. Plugging it into Equation (56) yields that $\text{border}_{\Pi}(\delta', \xi) = \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Border}_{\Pi}^{\delta', \xi} \right) \right] \leq m \cdot \sqrt{\xi}$, completing the proof of Proposition 4.19. \square

4.2.2 Attacking Approximated Pruned Protocols

Lemma 4.17 and Proposition 4.19 yield that for any protocol there is an (eventually polynomial-size) set of approximated pruned protocols, such that at least one of them has a small probability of visiting \mathbf{A} -controlled low-value transcripts. Hence, by Corollary 4.8, a (non-recursive) approximated biased-continuation attacker biases this approximated pruned protocol with similar success to that of the (ideal) biased-continuation attacker. A single recursion, however, is not guaranteed to be a good enough attacker. Hence, our next step is to argue the same for the *recursive* approximated biased-continuation attacker.

Unlike its non-recursive variant, the recursive approximated biased-continuation attacker might increase the probability of hitting low-value nodes. Let $\Pi = (\mathbf{A}, \mathbf{B})$ be a protocol in which the probability of hitting nodes whose value is smaller than δ (the set of δ -low nodes) is small. Consider the protocols $\Pi^{(k)} = (\mathbf{A}^{(k)}, \mathbf{B})$ and $\tilde{\Pi}^{(k)} = (\mathbf{A}^{(k, \xi, \delta)}, \mathbf{B})$, in which the recursive ideal and approximated biased-continuation attackers, respectively, take the role of \mathbf{A} . Since the probability of hitting δ -low nodes is small, the values of $\Pi^{(1)}$ and $\tilde{\Pi}^{(1)}$ are close (depending on the approximation guarantee ξ). For the next level of recursion, however, this might no longer be the case; in $\Pi^{(1)}$, every node has a higher value than in Π , so the δ -low set can only decrease. This is because the attacker is *always* successful in choosing a continuation leading to one (unless none exist). When replacing $\Pi^{(1)}$ with $\tilde{\Pi}^{(1)}$, however, the latter argument is no longer true. The approximated biased-continuator might fail to find a continuation leading to one, resulting in nodes whose value in $\tilde{\Pi}^{(1)}$ might be *smaller*

than in Π . Namely the δ -low value set might increase. Hence, it is no longer clear that the values of $\Pi^{(2)}$ and $\tilde{\Pi}^{(2)}$ are close.

Fortunately, as the next lemma shows, there is only a small probability that the situation described above will occur.

Lemma 4.21. *Let $0 < \delta \leq \delta' \leq \frac{1}{4}$, let $c = c(\delta)$ be according to Lemma 4.6, let $\xi \in (0, 1)$ and let $\tilde{\Pi} = (\tilde{\mathbf{A}}, \tilde{\mathbf{B}}) = \Pi^{[2\delta', \xi]}$ be the $(2\delta', \xi)$ -approximately pruned variant of a m -round protocol Π . Then*

$$\begin{aligned} & \text{SD} \left(\left\langle \mathbf{A}_{\tilde{\Pi}}^{(k)}, \tilde{\mathbf{B}} \right\rangle, \left\langle \mathbf{A}_{\tilde{\Pi}}^{(k, \xi, \delta')}, \tilde{\mathbf{B}} \right\rangle \right) \\ & \leq \phi^{\text{lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma) \\ & := k \cdot \frac{30^k \cdot m^k \cdot \prod_{i=1}^k \gamma_i}{\delta^{2k}} \cdot \left(\text{border}_{\Pi}(2\delta', \xi) + \frac{9 \cdot m \cdot \xi}{\delta'} \right) + \sum_{i=1}^k 2^{k-i+2} \cdot \frac{30^{k-i} \cdot m^{k-i} \cdot \prod_{j=i+1}^k \gamma_j}{\delta'^{2(k-i)} \cdot \gamma_i^c}, \end{aligned}$$

for every $k \in \mathbb{N}$ and $\gamma = (\gamma_1, \dots, \gamma_k)$ with $\gamma_i > 1$ for every $i \in [k]$.

As this lemma shows, there is only a small probability of hitting nodes whose value decreases when the recursive approximated biased-continuation attacker biases the approximated pruned protocol. This is because the value of each node can only increase when the recursive ideal biased-continuation attacker biases this protocol. The proof of Lemma 4.21 is given below, but we first use it to derive an important property of the recursive approximated biased-continuation attacker. In Section 4.1.1 (Proposition 4.9(2)) we showed that the *ideal* recursive biased-continuation attacker, when attacking the approximated pruned protocol, does not increase the probability of hitting any set of nodes by much. Using the above lemma, we can argue the same for the *approximated* recursive biased-continuation attacker when attacking the same protocol.

Proposition 4.22. *Let $0 < \delta \leq \delta' \leq \frac{1}{4}$, let $\xi \in (0, 1)$ and let $\tilde{\Pi} = (\tilde{\mathbf{A}}, \tilde{\mathbf{B}}) = \Pi^{[2\delta', \xi]}$ be the $(2\delta', \xi)$ -approximately pruned variant of an m -round protocol Π . Let \mathcal{F} be a frontier with $\Pr_{(\Pi)}[\text{desc}(\mathcal{F})] \leq \alpha$. Then*

$$\begin{aligned} \Pr_{\langle \mathbf{A}_{\tilde{\Pi}}^{(k, \delta', \xi)}, \tilde{\mathbf{B}} \rangle}[\text{desc}(\mathcal{F})] & \leq \phi^{\text{Bal}}(\alpha + 2 \cdot m \cdot \xi, \text{border}_{\Pi}(2\delta', \xi) + 4 \cdot m \cdot \xi / \delta', \delta, \gamma) \\ & \quad + \phi^{\text{lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma) \end{aligned}$$

for every $k \in \mathbb{N}$ and $\gamma = (\gamma_1, \dots, \gamma_k)$ with $\gamma_i > 1$ for every $i \in [k]$.⁴⁶

Proof. By Lemma 4.15, it holds that $\Pr_{\langle \tilde{\Pi} \rangle}[\text{desc}(\mathcal{F})] \leq \alpha + 2 \cdot m \cdot \xi$. By Lemma 4.17, it follows that

$$\Pr_{\langle \tilde{\Pi} \rangle} \left[\text{desc} \left(\text{Small}_{\tilde{\Pi}}^{\delta', \tilde{\mathbf{A}}} \right) \right] \leq \text{border}_{\Pi}(2\delta', \xi) + \frac{4 \cdot m \cdot \xi}{\delta'}.$$

Hence, Proposition 4.9(2) yields that

$$\Pr_{\langle \mathbf{A}_{\tilde{\Pi}}^{(k, \delta', \xi)}, \tilde{\mathbf{B}} \rangle}[\text{desc}(\mathcal{F})] \leq \phi^{\text{Bal}}(\alpha + 2 \cdot m \cdot \xi, \text{border}_{\Pi}(2\delta', \xi) + 4 \cdot m \cdot \xi / \delta', \delta, \gamma).$$

The proof of Proposition 4.22 now follows directly from Lemma 4.21. \square

⁴⁶See Proposition 4.9 and Lemma 4.21, for the definitions of ϕ^{Bal} and ϕ^{lt} respectively.

Finally, we analyze the efficiency (i.e., running time and number of randoms bits) of the approximated biased-continuation attacker, when attacking the approximated pruned protocols. This analysis assumes access to an honest-continuator and an estimator for the original protocol, and the efficiency is stated with respect to these algorithms.

Lemma 4.23. *Let $\delta \in (0, \frac{1}{4})$, let $\xi \in (0, \frac{1}{2})$, and let $\tilde{\Pi} = \Pi^{[2\delta, \xi]}$ be the $(2\delta, \xi)$ -approximately pruned variant of an m -round protocol Π . Assume the running times of $\text{Est}_{\tilde{\Pi}}^{\xi}$ and $\text{HonCont}_{\tilde{\Pi}}^{\xi}$ are T_{Est} and T_{HonCont} respectively, and that $\text{HonCont}_{\tilde{\Pi}}^{\xi}$ uses at most ρ_{HonCont} random bits. Then, for every $k \in \mathbb{N}$, algorithm $A_{\tilde{\Pi}}^{(k, \xi, \delta)}$ (see Algorithm 4.2) has the following properties.*

1. *It uses at most $k \cdot m^k \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^k \cdot \rho_{\text{HonCont}}$ random bits.*
2. *It runs in time (at most) $k \cdot m^k \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^k \cdot (T_{\text{Est}} + T_{\text{HonCont}})$.*

4.2.2.1 Proving Lemma 4.21

Proof of Lemma 4.21. Fix $\gamma_1, \dots, \gamma_k > 1$. The proof is by induction on k . For $k = 1$, plugging Lemma 4.17 into Corollary 4.8 yields that

$$\begin{aligned} \text{SD} \left(\left\langle A_{\tilde{\Pi}}^{(1)}, \tilde{\mathbf{B}} \right\rangle, \left\langle A_{\tilde{\Pi}}^{(1, \xi, \delta')}, \tilde{\mathbf{B}} \right\rangle \right) &\leq 2 \cdot m \cdot \gamma_1 \cdot \left(\xi + \text{border}_{\Pi}(2\delta', \xi) + \frac{4 \cdot m \cdot \xi}{\delta'} \right) + \frac{2}{\gamma_1^c} \\ &\leq \frac{30 \cdot m \cdot \gamma_1}{\delta'^2} \cdot \left(\text{border}_{\Pi}(2\delta', \xi) + \frac{9 \cdot m \cdot \xi}{\delta'} \right) + \frac{4}{\gamma_1^c}, \end{aligned}$$

as required.

Assume the lemma holds for $k - 1$; we prove it for k . The proof relies on the following lemma.

Lemma 4.24. *Let $\Pi = (A, B)$ and $\Pi' = (C, D)$ be two m -round protocols with the same control scheme, let $0 < \delta \leq \delta' \leq \frac{1}{4}$, let $c = c(\delta)$ be according to Lemma 4.6, and let $\xi \in (0, 1)$. Assume that*

1. $\chi_{\Pi} \equiv \chi_{\Pi'}$,
2. $\text{SD}(\langle \Pi \rangle, \langle \Pi' \rangle) \leq \alpha$, and
3. $\Pr_{\langle \Pi' \rangle} \left[\text{desc} \left(\text{Small}_{\Pi'}^{1.5\delta', C} \right) \right] \leq \beta$.

Then for every $\xi > 0$ and $\gamma > 1$, it holds that

$$\text{SD} \left(\left\langle A_{\tilde{\Pi}}^{(1, \xi, \delta')}, \tilde{\mathbf{B}} \right\rangle, \left\langle C_{\Pi'}^{(1)}, \tilde{\mathbf{D}} \right\rangle \right) \leq \frac{30 \cdot m \cdot \gamma}{\delta'^2} \cdot (\alpha + \xi + \beta) + \frac{4}{\gamma^c}.$$

The proof of Lemma 4.24 is given below, but first we use it to prove Lemma 4.21. For $i \in [k]$, let $\Pi_1^{(i)} = \left(A_{\tilde{\Pi}}^{(i)}, \tilde{\mathbf{B}} \right)$ and let $\Pi_2^{(i)} = \left(A_{\tilde{\Pi}}^{(i, \xi, \delta')}, \tilde{\mathbf{B}} \right)$. By this notation, we can write $\Pi_1^{(k)} = \left(A_{\Pi_1^{(k-1)}}^{(1)}, \tilde{\mathbf{B}} \right)$ and $\Pi_2^{(k)} = \left(A_{\Pi_2^{(k-1)}}^{(1, \xi, \delta')}, \tilde{\mathbf{B}} \right)$. Hence,

$$\text{SD} \left(\left\langle A_{\tilde{\Pi}}^{(k)}, \tilde{\mathbf{B}} \right\rangle, \left\langle A_{\tilde{\Pi}}^{(k, \xi, \delta')}, \tilde{\mathbf{B}} \right\rangle \right) = \text{SD} \left(\left\langle A_{\Pi_1^{(k-1)}}^{(1)}, \tilde{\mathbf{B}} \right\rangle, \left\langle A_{\Pi_2^{(k-1)}}^{(1, \xi, \delta')}, \tilde{\mathbf{B}} \right\rangle \right). \quad (57)$$

We would like to apply Lemma 4.24 with respect to $\Pi_1^{(k-1)}$ and $\Pi_2^{(k-1)}$. Indeed, these protocols share the same control scheme and common output function χ , and the induction hypothesis gives us a bound for $\text{SD}(\langle \Pi_1^{(k-1)} \rangle, \langle \Pi_2^{(k-1)} \rangle)$. It remains to bound $\Pr_{\langle \Pi_1^{(k-1)} \rangle} \left[\text{desc} \left(\text{Small}_{\Pi_1^{(k-1)}}^{1.5\delta', A} \right) \right]$. Note that the value of a node in $\Pi_1^{(i)}$ cannot be lower than its value in $\Pi_1^{(i-1)}$, and thus

$$\begin{aligned} \Pr_{\langle \Pi_1^{(k-1)} \rangle} \left[\text{desc} \left(\text{Small}_{\Pi_1^{(k-1)}}^{1.5\delta', A} \right) \right] &\leq \Pr_{\langle \Pi_1^{(k-1)} \rangle} \left[\text{desc} \left(\text{Small}_{\tilde{\Pi}}^{1.5\delta', A} \right) \right] \\ &\leq \Pr_{\langle \tilde{\Pi} \rangle} \left[\text{desc} \left(\text{Small}_{\tilde{\Pi}}^{1.5\delta', A} \right) \right] \cdot \prod_{i=1}^{k-1} \gamma_i + 2 \cdot \sum_{i=1}^{k-1} \frac{\prod_{j=i+1}^{k-1} \gamma_j}{\gamma_i^c} \\ &\leq \left(\text{border}_{\Pi}(2\delta', \xi) + \frac{8 \cdot m \cdot \xi}{\delta'} \right) \cdot \prod_{i=1}^{k-1} \gamma_i + 2 \cdot \sum_{i=1}^{k-1} \frac{\prod_{j=i+1}^{k-1} \gamma_j}{\gamma_i^c}. \end{aligned} \quad (58)$$

The second inequality follows Proposition 4.9(1) and the third one Lemma 4.17. By the induction hypothesis and Lemma 4.24, it holds that

$$\begin{aligned} &\text{SD} \left(\langle A_{\tilde{\Pi}}^{(k)}, \tilde{B} \rangle, \langle A_{\tilde{\Pi}}^{(k, \xi, \delta')}, \tilde{B} \rangle \right) \\ &\leq \frac{30 \cdot m \cdot \gamma_k}{\delta'^2} \cdot \left((k-1) \cdot \frac{30^{k-1} \cdot m^{k-1} \cdot \prod_{i=1}^{k-1} \gamma_i}{\delta'^{2(k-1)}} \cdot \left(\text{border}_{\Pi}(2\delta', \xi) + \frac{9 \cdot m \cdot \xi}{\delta'} \right) \right. \\ &\quad \left. + \sum_{i=1}^{k-1} 2^{k-i+1} \cdot \frac{30^{k-1-i} \cdot m^{k-1-i} \cdot \prod_{j=i+1}^{k-1} \gamma_j}{\delta'^{2(k-1-i)} \cdot \gamma_i^c} \right. \\ &\quad \left. + \left(\text{border}_{\Pi}(2\delta', \xi) + \frac{8 \cdot m \cdot \xi}{\delta'} \right) \cdot \prod_{i=1}^{k-1} \gamma_i + 2 \cdot \sum_{i=1}^{k-1} \frac{\prod_{j=i+1}^{k-1} \gamma_j}{\gamma_i^c} + \xi \right) + \frac{4}{\gamma_k^c} \\ &= \frac{30 \cdot m \cdot \gamma_k}{\delta'^2} \cdot \left((k-1) \cdot \frac{30^{k-1} \cdot m^{k-1} \cdot \prod_{i=1}^{k-1} \gamma_i}{\delta'^{2(k-1)}} \cdot \left(\text{border}_{\Pi}(2\delta', \xi) + \frac{9 \cdot m \cdot \xi}{\delta'} \right) \right. \\ &\quad \left. + \left(\text{border}_{\Pi}(2\delta', \xi) + \frac{8 \cdot m \cdot \xi}{\delta'} \right) \cdot \prod_{i=1}^{k-1} \gamma_i + \xi \right) \\ &\quad + \frac{30 \cdot m \cdot \gamma_k}{\delta'^2} \cdot \left(\sum_{i=1}^{k-1} 2^{k-i+1} \cdot \frac{30^{k-1-i} \cdot m^{k-1-i} \cdot \prod_{j=i+1}^{k-1} \gamma_j}{\delta'^{2(k-1-i)} \cdot \gamma_i^c} + 2 \cdot \sum_{i=1}^{k-1} \frac{\prod_{j=i+1}^{k-1} \gamma_j}{\gamma_i^c} \right) + \frac{4}{\gamma_k^c}. \end{aligned}$$

The induction proof now follows by grouping together the summands in the parentheses. This concludes the proof of Lemma 4.21. \square

Proof of Lemma 4.24. To prove the lemma, we make use of the biased-continuation attacker being robust.

Lemma 4.25 (robustness lemma). *Let $\Pi = (A, B)$ and $\Pi' = (C, D)$ be two m -round protocols, let $\delta \in (0, \frac{1}{2}]$, and let $c = c(\delta)$ according to Lemma 4.6. Assume that $\text{SD}(\langle \Pi \rangle, \langle \Pi' \rangle) \leq \alpha$, that $\chi_{\Pi} \equiv \chi_{\Pi'}$, and that Π and Π' have the same control scheme. Then*

$$\text{SD} \left(\langle A_{\Pi}^{(1)}, B \rangle, \langle C_{\Pi'}^{(1)}, D \rangle \right) \leq \frac{3 \cdot m \cdot \gamma}{\delta'} \cdot \left(\alpha + \Pr_{(A, B)} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C} \right) \right] \right) + \frac{2}{\gamma^c},$$

for every $\delta' \geq \delta$ and $\gamma \geq 1$, where $A^{(1)}$ and $C^{(1)}$ are as in Algorithm 3.2.

Namely, the biased-continuation attacker does not make similar protocols too dissimilar. To get an intuition about this robustness property, recall that the biased-continuation attacker, when attacking Π , chooses a random 1-leaf according to $\langle \Pi \rangle$, the leaf distribution of Π . Since Π and Π' are similar, however, the attacker can instead sample from $\langle \Pi' \rangle$, while making similar decisions throughout its operation. So, the biased-continuation attacker is robust to the distribution from which it samples. Lemma 4.25 is proven below.

We now return to proving Lemma 4.24. The proof proceeds in two steps: first, we apply Lemma 4.25 (robustness lemma) to show that after the (ideal) biased-continuation attacker takes the role of the left-hand party in Π and Π' , the leaf distributions of these protocols remain close; second, we apply Corollary 4.8 (ideal-to-approximated biased-continuation attacker) to show that replacing the attacker of the left-hand party in Π with its approximated variant, the leaf distributions of these protocols remain close.

In order to apply Lemma 4.25, we first need to bound $\Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C} \right) \right]$. Let $\mathcal{F} = \text{frnt} \left(\text{Small}_{\Pi}^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C} \right)$, let $\mathcal{F}_1 = \{u \in \mathcal{F} : \text{val}((\Pi')_u) \geq 1.5\delta'\}$, and let $\mathcal{F}_2 = \{u \in \mathcal{F} : \text{val}((\Pi)_u) < 1.5\delta'\}$. Since $\mathcal{F} \subseteq \mathcal{F}_1 \cup \mathcal{F}_2$, it suffices to bound $\Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}_1)]$ and $\Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}_2)]$, which we do separately.

Bounding \mathcal{F}_1 : Nodes in \mathcal{F}_1 must have small value in Π but large value in Π' . Since $\langle \Pi \rangle$ and $\langle \Pi' \rangle$ are close, the probability of reaching such nodes is small.

Let $\mu = \Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}_1)]$, $\mu' = \Pr_{\langle \Pi' \rangle} [\text{desc}(\mathcal{F}_1)]$. By noting that every node in \mathcal{F}_1 must belong to $\text{Small}_{\Pi}^{\delta', A}$, it follows that $\Pr_{\langle \Pi \rangle} [\mathcal{L}_1(\Pi) \mid \text{desc}(\mathcal{F}_1)] \leq \delta'$. The assumption (1) of the lemma and the definition of \mathcal{F}_1 yield, however, that $\Pr_{\langle \Pi' \rangle} [\mathcal{L}(\Pi) \mid \text{desc}(\mathcal{F}_1)] \geq 1.5\delta'$. By Proposition 4.18 it follows that

$$\Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}_1)] \leq \alpha \cdot \frac{1 + 1.5\delta'}{0.5\delta'} \leq \frac{4\alpha}{\delta'}.$$

The last inequality holds since, by assumption, $\delta' \leq 1/4$.

Bounding \mathcal{F}_2 : The definition of \mathcal{F}_2 , the assumption that Π and Π' have the same control scheme, and assumption (3), yield that $\Pr_{\langle \Pi' \rangle} [\text{desc}(\mathcal{F}_2)] \leq \beta$. Hence, the assumption that $\text{SD}(\langle \Pi \rangle, \langle \Pi' \rangle) \leq \alpha$ (assumption (2) of the lemma) yields that $\Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}_2)] \leq \alpha + \beta$.

Combining the two bounds, it follows that $\Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C} \right) \right] \leq 5\alpha/\delta' + \beta$. We can apply Lemma 4.25 and derive

$$\text{SD} \left(\langle A_{\Pi}^{(1)}, B \rangle, \langle C_{\Pi'}^{(1)}, D \rangle \right) \leq \frac{3 \cdot m \cdot \gamma}{\delta'} \cdot \left(\alpha + \frac{5\alpha}{\delta'} + \beta \right) + \frac{2}{\gamma^c}. \quad (59)$$

Our next step is to apply Corollary 4.8. To do so we need to bound $\Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right]$, but since it is clear that $\Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right] \leq \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C} \right) \right]$, it follows that $\Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \right) \right] \leq 5\alpha/\delta' + \beta$. Applying Corollary 4.8, we derive

$$\text{SD} \left(\langle A_{\Pi}^{(1)}, B \rangle, \langle A_{\Pi}^{(1, \xi, \delta')}, B \rangle \right) \leq 2 \cdot m \cdot \gamma \cdot \left(\xi + \frac{5\alpha}{\delta'} + \beta \right) + \frac{2}{\gamma^c}. \quad (60)$$

Finally, combining the last two inequalities and using the triangle inequality completes the proof of Lemma 4.24. \square

Proving The Robustness Lemma — Lemma 4.25

Proof of Lemma 4.25. We use Lemma 2.14. Define the random function $f: \mathcal{V}(\Pi) \mapsto \mathcal{L}(\Pi)$ as follows: given $u \in \mathcal{V}(\Pi)$, if **A** controls u return $\ell \leftarrow \langle \Pi_u \rangle$ such that $\chi_\Pi(\ell) = 1$, and otherwise, i.e., if **B** controls u , return $\ell \leftarrow \langle \Pi_u \rangle$. The random function $g: \mathcal{V}(\Pi) \mapsto \mathcal{L}(\Pi)$ is analogously defined with respect to protocol Π' .⁴⁷ For function **O** with range in $\mathcal{L}(\Pi)$, let H^O be the following algorithm:

Algorithm 4.26 (H).

Oracle: **O**.

State: node u , set to λ at the start of the execution.

Operation:

1. Repeat for m times:
 - (a) Set $\ell := O(u)$.
 - (b) Set $u := u \circ \ell_i$, where i is the current iteration.
2. Output u .

It is easy to verify that $H^f \equiv \langle A_\Pi^{(1)}, B \rangle$ and $H^g \equiv \langle C_{\Pi'}^{(1)}, D \rangle$. Hence, it suffices to upper-bound $SD(E^f, E^g)$. For $i \in [m]$, let D_i to be i 'th node in a random execution of Π (such a node consists of $i - 1$ bits). We use the next claim, proven below.

Claim 4.27. $E_{u \leftarrow D_i} [SD(f(u), g(u))] \leq \frac{2\alpha}{\delta^r} + \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_\Pi^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C} \right) \right]$.

Let Q_i denote the i 'th query to f in a random execution of H^f (note that by construction, such a query always exists) and let $Q = (Q_1, \dots, Q_m)$. By construction, for $u \in \mathcal{V}(\Pi)$ with $|u| = i - 1$, $Q_i(u)$ is the probability that u is visited in a random execution of $\langle A_\Pi^{(1)}, B \rangle$. We get

$$\begin{aligned} \Pr_{(q_1, \dots, q_m) \leftarrow Q} [\exists i \in [m]: q_i \neq \perp \wedge Q_i(q_i) > \gamma \cdot D_i(q_i)] &= \Pr_{\langle A^{(1)}, B \rangle} [\text{desc}(\text{UnBal}_\Pi^\gamma)] \\ &\leq \gamma \cdot \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_\Pi^{\delta', A} \right) \right] + \frac{2}{\gamma^c}, \end{aligned}$$

where the inequality follows from Corollary 4.7.

The proof of Lemma 4.25 now follows by Lemma 2.14, letting $k = m$, $a = \frac{2\alpha}{\delta^r} + \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_\Pi^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C} \right) \right]$, $\lambda = \gamma$ and $b = \gamma \cdot \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_\Pi^{\delta', A} \right) \right] + \frac{2}{\gamma^c}$. \square

⁴⁷The sets $\mathcal{V}(\Pi)$ and $\mathcal{V}(\Pi')$, as well as the sets $\mathcal{L}(\Pi)$ and $\mathcal{L}(\Pi')$, are identical, as the both describe nodes in the complete binary tree of height m . See Section 2 for further details.

Proof of Claim 4.27. Let $\mathcal{V}_i(\Pi) = \{v \in \mathcal{V}(\Pi) : |v| = i - 1\}$, $\mathcal{V}_i^A(\Pi) = \mathcal{V}_i(\Pi) \cap \text{Ctrl}_\Pi^A$ and $\mathcal{V}_i^B(\Pi) = \mathcal{V}_i(\Pi) \cap \text{Ctrl}_\Pi^B$. Compute

$$\begin{aligned} \mathbb{E}_{u \leftarrow D_i} [\text{SD}(f(u), g(u))] &= \sum_{u \in \mathcal{V}_i(\Pi)} D_i(u) \cdot \text{SD}(f(u), g(u)) \\ &= \sum_{u \in \mathcal{V}_i^A(\Pi)} D_i(u) \cdot \text{SD}(f(u), g(u)) + \sum_{u \in \mathcal{V}_i^B(\Pi)} D_i(u) \cdot \text{SD}(f(u), g(u)). \end{aligned} \quad (61)$$

In the rest of the proof we show that

$$\begin{aligned} \sum_{u \in \mathcal{V}_i^A(\Pi)} D_i(u) \cdot \text{SD}(f(u), g(u)) &\leq \frac{1}{\delta'} \cdot \sum_{u \in \mathcal{V}_i^A(\Pi)} D_i(u) \cdot \text{SD}(\langle \Pi_u \rangle, \langle \Pi'_u \rangle) \\ &\quad + \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C} \right) \right], \end{aligned} \quad (62)$$

that

$$\sum_{u \in \mathcal{V}_i^B(\Pi)} D_i(u) \cdot \text{SD}(f(u), g(u)) \leq \sum_{u \in \mathcal{V}_i^B(\Pi)} D_i(u) \cdot \text{SD}(\langle \Pi_u \rangle, \langle \Pi'_u \rangle), \quad (63)$$

and that

$$\sum_{u \in \mathcal{V}_i(\Pi)} D_i(u) \cdot \text{SD}(\langle \Pi_u \rangle, \langle \Pi'_u \rangle) \leq 2 \cdot \text{SD}(\langle \Pi \rangle, \langle \Pi' \rangle). \quad (64)$$

Plugging Equations (62) to (64) into Equation (61) completes the proof Claim 4.27.

Proof of Equation (62): Let $u \in \mathcal{V}_i^A(\Pi)$. By the definition of f , and since u is under A 's control, it follows that $\Pr[f(u) = \ell] = \langle \Pi_u \rangle(\ell) / \text{val}(\Pi_u)$ if $\chi_\Pi(\ell) = 1$, and $\Pr[f(u) = \ell] = 0$ otherwise. Since Π and Π' have the same control scheme, the same holds for $g(u)$ with respect to Π' . Let $\mathcal{S}'_u \subseteq \mathcal{L}_1(\Pi)$ be the set with $\text{SD}(f(u), g(u)) = \sum_{\ell \in \mathcal{S}'_u} (\Pr[f(u) = \ell] - \Pr[g(u) = \ell]) = \sum_{\ell \in \mathcal{L}_1(\Pi) \setminus \mathcal{S}'_u} (\Pr[g(u) = \ell] - \Pr[f(u) = \ell])$.⁴⁸ Define $\mathcal{S}_u \subseteq \mathcal{L}_1(\Pi)$ as follows: if $\text{val}(\Pi_u) \geq \text{val}(\Pi'_u)$ let $\mathcal{S}_u = \mathcal{S}'_u$; otherwise let $\mathcal{S}_u = \mathcal{L}_1(\Pi) \setminus \mathcal{S}'_u$. It follows that

$$\begin{aligned} \sum_{u \in \mathcal{V}_i^A(\Pi)} D_i(u) \cdot \text{SD}(f(u), g(u)) &\leq \sum_{\substack{u \in \mathcal{V}_i^A(\Pi) : \\ \text{val}(\Pi_u) \geq \text{val}(\Pi'_u) \geq \delta'}} D_i(u) \cdot \sum_{\ell \in \mathcal{S}_u} \left(\frac{\langle \Pi_u \rangle(\ell)}{\text{val}(\Pi_u)} - \frac{\langle \Pi'_u \rangle(\ell)}{\text{val}(\Pi'_u)} \right) \\ &\quad + \sum_{\substack{u \in \mathcal{V}_i^A(\Pi) : \\ \text{val}(\Pi'_u) > \text{val}(\Pi_u) \geq \delta'}} D_i(u) \cdot \sum_{\ell \in \mathcal{S}_u} \left(\frac{\langle \Pi'_u \rangle(\ell)}{\text{val}(\Pi'_u)} - \frac{\langle \Pi_u \rangle(\ell)}{\text{val}(\Pi_u)} \right) \\ &\quad + \sum_{\substack{u \in \mathcal{V}_i^A(\Pi) : \\ \text{val}(\Pi_u) < \delta' \vee \text{val}(\Pi'_u) < \delta'}} D_i(u). \end{aligned} \quad (65)$$

Assume that $\text{val}(\Pi_u) \geq \text{val}(\Pi'_u)$. The definition of \mathcal{S}_u implies that $\langle \Pi_u \rangle(\ell) / \text{val}(\Pi_u) \geq \langle \Pi'_u \rangle(\ell) / \text{val}(\Pi'_u)$ for every $\ell \in \mathcal{S}_u$. But since $\text{val}(\Pi_u) / \text{val}(\Pi'_u) \geq 1$, the latter yields that

⁴⁸Note that it must be the case that $\mathcal{S}'_u \subseteq \mathcal{L}_1(\Pi)$, since $\Pr[f(u) = \ell] = \Pr[g(u) = \ell] = 0$, for every ℓ with $\chi_\Pi(\ell) = 0$, which follows from the assumption that $\chi_\Pi \equiv \chi_{\Pi'}$.

$\langle \Pi_u \rangle(\ell) \geq \langle \Pi'_u \rangle(\ell)$ for every $\ell \in \mathcal{S}_u$. Using this observation, we bound the first summand in the right-hand side of Equation (65).

$$\begin{aligned}
& \sum_{\substack{u \in \mathcal{V}_i^A(\Pi): \\ \text{val}(\Pi_u) \geq \text{val}(\Pi'_u) \geq \delta'}} D_i(u) \cdot \sum_{\ell \in \mathcal{S}_u} \left(\frac{\langle \Pi_u \rangle(\ell)}{\text{val}(\Pi_u)} - \frac{\langle \Pi'_u \rangle(\ell)}{\text{val}(\Pi'_u)} \right) \\
& \leq \sum_{\substack{u \in \mathcal{V}_i^A(\Pi): \\ \text{val}(\Pi_u) \geq \text{val}(\Pi'_u) \geq \delta'}} \frac{D_i(u)}{\text{val}(\Pi')} \cdot \sum_{\ell \in \mathcal{S}_u} (\langle \Pi_u \rangle(\ell) - \langle \Pi'_u \rangle(\ell)) \\
& \leq \frac{1}{\delta'} \sum_{\substack{u \in \mathcal{V}_i^A(\Pi): \\ \text{val}(\Pi_u) \geq \text{val}(\Pi'_u) \geq \delta'}} D_i(u) \cdot \sum_{\ell \in \mathcal{S}_u} (\langle \Pi_u \rangle(\ell) - \langle \Pi'_u \rangle(\ell)) \\
& \leq \frac{1}{\delta'} \sum_{\substack{u \in \mathcal{V}_i^A(\Pi): \\ \text{val}(\Pi_u) \geq \text{val}(\Pi'_u) \geq \delta'}} D_i(u) \cdot \text{SD}(\langle \Pi_u \rangle, \langle \Pi'_u \rangle),
\end{aligned} \tag{66}$$

where the second inequality follows since $\sum_{\ell \in \mathcal{S}_u} (\langle \Pi_u \rangle(\ell) - \langle \Pi'_u \rangle(\ell)) \geq 0$, as argued above. Similar calculations, and using the symmetry of statistical distance, we bound the second summand in the right-hand side of Equation (65):

$$\sum_{\substack{u \in \mathcal{V}_i^A(\Pi): \\ \text{val}(\Pi'_u) \geq \text{val}(\Pi_u) \geq \delta'}} D_i(u) \cdot \sum_{\ell \in \mathcal{S}_u} \left(\frac{\langle \Pi'_u \rangle(\ell)}{\text{val}(\Pi'_u)} - \frac{\langle \Pi_u \rangle(\ell)}{\text{val}(\Pi_u)} \right) \leq \frac{1}{\delta'} \sum_{\substack{u \in \mathcal{V}_i^A(\Pi): \\ \text{val}(\Pi'_u) \geq \text{val}(\Pi_u) \geq \delta'}} D_i(u) \cdot \text{SD}(\langle \Pi_u \rangle, \langle \Pi'_u \rangle). \tag{67}$$

Finally, to bound the third summand in the right-hand side of Equation (65), we note that it sums over (not all) $u \in \text{Small}_{\Pi}^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C}$. Since D_i simply samples a random partial transcript from Π , we derive the following bound:

$$\sum_{\substack{u \in \mathcal{V}_i^A(\Pi): \\ \text{val}(\Pi_u) < \delta' \vee \text{val}(\Pi'_u) < \delta'}} D_i(u) \leq \Pr_{\langle \Pi \rangle} \left[\text{desc} \left(\text{Small}_{\Pi}^{\delta', A} \cup \text{Small}_{\Pi'}^{\delta', C} \right) \right]. \tag{68}$$

Plugging Equations (66) to (68) into Equation (65) yields Equation (62).

Proof of Equation (63): Since it is the right-hand party who controls u in Π and in Π' , it follows that $\text{SD}(f(u), g(u)) = \text{SD}(\langle \Pi_u \rangle, \langle \Pi'_u \rangle)$. Equation (63) follows immediately.

Proof of Equation (64): Using the definition of D_i , we can write

$$\begin{aligned}
\sum_{u \in \mathcal{V}_i(\Pi)} D_i(u) \cdot \text{SD}(\langle \Pi_u \rangle, \langle \Pi'_u \rangle) &= \sum_{u \in \mathcal{V}_i(\Pi)} \nu_{\Pi}(u) \cdot \frac{1}{2} \sum_{\ell \in \mathcal{L}(\Pi_u)} |\nu_{\Pi_u}(\ell) - \nu_{\Pi'_u}(\ell)| \\
&= \frac{1}{2} \sum_{\ell \in \mathcal{L}(\Pi)} \left| \nu_{\Pi}(\ell_{1, \dots, i-1}) \cdot \nu_{\Pi_{\ell_1, \dots, i-1}}(\ell) - \nu_{\Pi}(\ell_{1, \dots, i-1}) \cdot \nu_{\Pi'_{\ell_1, \dots, i-1}}(\ell) \right| \\
&= \text{SD}(\langle \Pi \rangle, \langle \Pi'' \rangle),
\end{aligned}$$

for $\langle \Pi'' \rangle(\ell) := \nu_{\Pi}(\ell_{1,\dots,i-1}) \cdot \nu_{\Pi'_{\ell_1,\dots,i-1}}(\ell)$.

We prove that $\text{SD}(\langle \Pi' \rangle, \langle \Pi'' \rangle) \leq \text{SD}(\langle \Pi' \rangle, \langle \Pi \rangle)$, and Equation (64) follows from the triangle inequality. Let h be the random function that, given $\ell \in \mathcal{L}(\Pi)$, returns $\ell' \leftarrow \langle \Pi'_{\ell_1,\dots,i-1} \rangle$. Therefore, $h(\langle \Pi' \rangle) \equiv \langle \Pi' \rangle$ and $h(\langle \Pi \rangle) \equiv \langle \Pi'' \rangle$, and this completes the proof.

This completes the proof of Equations (62) to (64), and thus the proof of Claim 4.27. \square

4.2.2.2 Proving Lemma 4.23

The attacker $A_{\tilde{\Pi}}^{(k,\xi,\delta)}$ requires a (ξ, δ) -biased continuator for the protocol $(A_{\tilde{\Pi}}^{(k-1,\xi,\delta)}, \tilde{B})$. Since the efficiency of $A_{\tilde{\Pi}}^{(k,\xi,\delta)}$ is simply that of calling this continuator, we need to show how to implement this continuator and analyze its efficiency. This implementation relies on the fact that the protocol $(A_{\tilde{\Pi}}^{(i,\xi,\delta)}, \tilde{B})$ is *stateless* for every $0 \leq i \leq k-1$ (for the case $i=0$, we simply get the approximated pruned protocol of Π , which is by definition stateless). We start by showing how to implement an *honest* continuator for a stateless protocol, and then show how a *biased* continuator can be implemented using an honest continuator.

Honest continuator for stateless protocols. For stateless protocols (i.e., the parties maintain no state), providing (perfect) honest continuation is immediate.

Algorithm 4.28 (HonContSL_{Π}).

Input: transcript $u \in \{0, 1\}^*$.

Operation:

1. Set $t := u$.
2. Repeat until $t \in \mathcal{L}(\Pi)$:
 - (a) Let C be the party that controls t .
 - (b) Choose uniformly at random coins r_C for this round.
 - (c) Set $t := t \circ C(t; r_C)$.
3. Return t .

Claim 4.29. Assume that Π is stateless. Then HonContSL_{Π} of Algorithm 4.28 is a 0-honest continuator for Π .

Proof. Immediate. \square

From honest continuation to biased continuation. Turning an honest continuator into a biased continuator is also straightforward: given a transcript u and a bit b toward which the continuator should bias, sample sufficiently many honest continuations for u , and return the first continuation whose common output is b . Indeed, if the transcript value (i.e., expected outcome) is close enough to b , then with high probability the above process indeed returns a biased continuation.

Algorithm 4.30 ($\text{BiasedCont}_{\Pi}^{(\xi,\delta,\text{HC})}$).

Parameters: $\xi, \delta \in (0, 1)$.

Oracle: algorithm HC.

Input: $u \in \mathcal{V}(\Pi)$ and $b \in \{0, 1\}$.

Operation:

1. For $i = 1$ to $\left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil$:
 - (a) Set $\ell := \text{HC}(u)$.
 - (b) If $\chi_\Pi(\ell) = b$, return $\ell_{|u|+1}$.
2. Return \perp .

Claim 4.31. Let Π be an m -round protocol, let $\xi, \xi', \delta \in (0, 1)$ and let $t = \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil$. Assume that HC is a ξ' -honest continuator for Π . Then $\text{BiasedCont}_\Pi^{(\xi, \delta, \text{HC})}$ is a $((t+1) \cdot \xi' + \xi, \delta)$ -biased continuator for Π .

Proof. We show that for every $u \in \mathcal{V}(\Pi)$ with $\text{SD}(\text{HC}(u), \text{HonCont}(u)) \leq \xi'$ and $\text{val}(\Pi_u) \geq \delta$, it holds that

$$\text{SD}\left(\text{BiasedCont}_\Pi^{(\xi, \delta, \text{HC})}(u, 1), \text{BiasedCont}(u, 1)\right) \leq t \cdot \xi' + \xi. \quad (69)$$

This suffices to complete the proof since that case for $\text{val}(\Pi_u) \leq 1 - \delta$ is analogous and since the probability that Π generates a node u such that $\text{SD}(\text{HC}(u), \text{HonCont}_\Pi(u)) > \xi'$ is at most ξ' .

Let $u \in \mathcal{V}(\Pi)$ with $\text{SD}(\text{HC}(u), \text{HonCont}_\Pi(u)) \leq \xi'$ and $\text{val}(\Pi_u) \geq \delta$. Define the following algorithm, implementing the well-known *rejection sampling* strategy.

Algorithm 4.32 ($\widehat{\text{BiasedCont}}$).

Operation:

1. Do (forever):
 - (a) Set $\ell := \text{HonCont}(u)$.
 - (b) If $\chi_\Pi(\ell) = 1$, return $\ell_{|u|+1}$.

It is not difficult to verify that the probability that $\widehat{\text{BiasedCont}}$ does not halt is zero, and thus $\widehat{\text{BiasedCont}} \equiv \text{BiasedCont}(u, 1)$, and that

$$\text{SD}\left(\text{BiasedCont}_\Pi^{(\xi, \delta, \text{HonCont})}(u, 1), \widehat{\text{BiasedCont}}\right) \leq \Pr\left[\text{BiasedCont}_\Pi^{(\xi, \delta, \text{HonCont})}(u, 1) = \perp\right]. \quad (70)$$

We compute

$$\begin{aligned} \Pr\left[\text{BiasedCont}_\Pi^{(\xi, \delta, \text{HonCont})}(u, 1) = \perp\right] &= \left(\Pr_{\ell \leftarrow \text{HonCont}(u)}[\chi_\Pi(\ell) = 0]\right)^t \\ &\leq (1 - \delta)^t \\ &\leq \xi, \end{aligned}$$

where the first inequality follows since $\text{val}(\Pi_u) \geq \delta$ and the last inequality follows from the choice of t . Moreover, since $\text{BiasedCont}_{\Pi}^{(\xi, \delta, \text{HC})}$ makes t calls to its oracle, the assumption that $\text{SD}(\text{HonCont}(u), \text{HC}(u)) \leq \xi^t$ and a standard hybrid argument yield that

$$\text{SD}\left(\text{BiasedCont}_{\Pi}^{(\xi, \delta, \text{HonCont})}(u, 1), \text{BiasedCont}_{\Pi}^{(\xi, \delta, \text{HC})}(u, 1)\right) \leq t \cdot \xi,$$

which completes the proof. \square

Having developed the necessary tools, we can now analyze the efficiency of the pruning-in-the-head attacker.

Proof of Lemma 4.23. We begin by giving a proof for Item 1 by showing that $A_{\Pi}^{(k, \xi, \delta)}$ uses at most $\sum_{i=1}^k m^i \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^i \cdot \rho_{\text{HonCont}}$ random bits. We then explain how to extend this analysis to prove Item 2.

For the base case $k = 1$, note that by definition, the protocol $\tilde{\Pi}$ is stateless. Thus, by Claim 4.29, algorithm $\text{HonContSL}_{\tilde{\Pi}}$ is a 0-honest continuator for $\tilde{\Pi}$, which uses at most $m \cdot \rho_{\text{HonCont}}$ random coins (in $\tilde{\Pi}$ the parties make a single call to $\text{Est}_{\tilde{\Pi}}^{\xi}$ and $\text{HonCont}_{\tilde{\Pi}}^{\xi}$ for every round; $\text{Est}_{\tilde{\Pi}}^{\xi}$ is assumed to be deterministic and thus no random coins are needed for calling it). Note that a call to $A_{\tilde{\Pi}}^{(1, \xi, \delta)}$ simply calls $\text{BiasedCont}_{\tilde{\Pi}}^{\xi, \delta}$, the fixed (ξ, δ) -biased continuator for $\tilde{\Pi}$ (see Definition 4.1). We now set $\text{BiasedCont}_{\tilde{\Pi}}^{\xi, \delta} := \text{BiasedCont}_{\tilde{\Pi}}^{(\xi, \delta, \text{HonContSL}_{\tilde{\Pi}})}$, which, by Claim 4.31 and since $\text{HonContSL}_{\tilde{\Pi}}$ is a 0-honest continuator for $\tilde{\Pi}$, is a (ξ, δ) -biased continuator for $\tilde{\Pi}$.⁴⁹ By the definition of $\text{BiasedCont}_{\tilde{\Pi}}^{(\xi, \delta, \text{HonContSL}_{\tilde{\Pi}})}$ (in Algorithm 4.30), it makes $\left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil$ calls to $\text{HonContSL}_{\tilde{\Pi}}$, each with fresh randomness. All in all, $A_{\tilde{\Pi}}^{(1, \xi, \delta)}$ uses at most $m \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil \cdot \rho_{\text{HonCont}}$ random bits, which completes the proof of the base case.

Assume the lemma holds for $k - 1$, and let $\tilde{\Pi}^{(k-1)} = (A_{\Pi}^{(k-1, \xi, \delta)}, \tilde{B})$, where \tilde{B} is the party taking the role of B in $\tilde{\Pi}$. As in the base case, note that $A_{\tilde{\Pi}}^{(k, \xi, \delta)}$ simply calls $\text{BiasedCont}_{\tilde{\Pi}^{(k-1)}}^{\xi, \delta}$, the fixed (ξ, δ) -biased continuator for $\tilde{\Pi}^{(k-1)}$. We now set $\text{BiasedCont}_{\tilde{\Pi}^{(k-1)}}^{\xi, \delta} := \text{BiasedCont}_{\tilde{\Pi}^{(k-1)}}^{(\xi, \delta, \text{HonContSL}_{\tilde{\Pi}^{(k-1)}})}$ and argue that this is valid (i.e., the latter algorithm is a (ξ, δ) -biased continuator for $\tilde{\Pi}^{(k-1)}$).

By definition, $A_{\tilde{\Pi}}^{(k-1, \xi, \delta)}$ is stateless, and thus $\tilde{\Pi}^{(k-1)}$ is a stateless protocol. By Claim 4.29, algorithm $\text{HonContSL}_{\tilde{\Pi}^{(k-1)}}$ is a 0-honest continuator for $\tilde{\Pi}^{(k-1)}$. Furthermore, for every A 's turn, $\text{HonContSL}_{\tilde{\Pi}^{(k-1)}}$ chooses random bits for $A_{\tilde{\Pi}}^{(k-1, \xi, \delta)}$, which by the induction hypothesis are at most $\sum_{i=1}^{k-1} m^i \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^i \cdot \rho_{\text{HonCont}}$; and for every B 's turn $\text{HonContSL}_{\tilde{\Pi}^{(k-1)}}$ chooses random bits for \tilde{B} , which are at most ρ_{HonCont} . Overall, $\text{HonContSL}_{\tilde{\Pi}^{(k-1)}}$ uses at most $m \cdot \sum_{i=1}^{k-1} m^i \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^i \cdot \rho_{\text{HonCont}} + m \cdot \rho_{\text{HonCont}}$ random bits. Now, same arguments from the base case yield that $\text{BiasedCont}_{\tilde{\Pi}^{(k-1)}}^{(\xi, \delta, \text{HonContSL}_{\tilde{\Pi}^{(k-1)}})}$ is a (ξ, δ) -biased continuator for $\tilde{\Pi}^{(k-1)}$ that uses

⁴⁹All the previously claimed properties of $A_{\Pi}^{(i, \xi, \delta)}$ are true for an arbitrary fixing of a (ξ, δ) -biased continuator. Specifically, they are true for this fixing.

at most

$$\begin{aligned} & \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil \cdot \left(m \cdot \sum_{i=1}^{k-1} m^i \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^i \cdot \rho_{\widetilde{\text{HonCont}}} + m \cdot \rho_{\widetilde{\text{HonCont}}} \right) = \\ & \sum_{i=1}^k m^i \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^i \cdot \rho_{\widetilde{\text{HonCont}}}, \end{aligned}$$

random bits, which completes the induction step.

To see that Item 2 holds, note that the running time of $\text{HonContSL}_{\widetilde{\Pi}}$ is $m \cdot (T_{\text{Est}} + T_{\text{HonCont}})$. The same calculations from above can be used to prove this item by replacing ρ_{HonCont} with $(T_{\text{Est}} + T_{\text{HonCont}})$. \square

4.3 The Pruning-in-the-Head Attacker

In the previous section we showed that the recursive approximated biased-continuation attacker successfully biases protocols, as long as these protocols are close to being pruned. Most protocols, however, do not possess the latter property. In this section we design an attacker that biases *any* protocol, while relying on the results of the previous section. This attacker applies the approximated biased-continuation attacker as if the attacked protocol is pruned, until it reaches a low or high value node, and then it switches its behavior to act honestly.

We begin with providing an intuition as to why this approach works in the ideal case. Consider the ideal pruned variant of a protocol pruned at some constant δ , which we denote by $\Pi^{[\delta]} = (\mathbf{A}^{[\delta]}, \mathbf{B}^{[\delta]})$ (see Definition 4.10). The ideal pruned protocol is itself a protocol, and the results of Section 3 apply to it. Specifically, and without loss of generality, Theorem 3.3 yields that $(\mathbf{A}^{[\delta]})^{(k)}$ successfully biases $\Pi^{[\delta]}$. Let $\mathbf{A}^{(k,\delta)}$ be the following attacker: until reaching a pruned node according to δ (i.e., a node whose value is lower than δ or higher than $1-\delta$), it acts like $(\mathbf{A}^{[\delta]})^{(k)}$; when reaching a pruned node, and in the rest of the execution, it acts like the honest party \mathbf{A} . $\mathbf{A}^{(k,\delta)}$ “thinks” — “in its head” — it is actually attacking the pruned variant of the protocol, instead of the original protocol.

We argue that $\mathbf{A}^{(k,\delta)}$ biases the original protocol almost as well as $(\mathbf{A}^{[\delta]})^{(k)}$ biases the ideal pruned protocol. Consider the protocols $((\mathbf{A}^{[\delta]})^{(k)}, \mathbf{B}^{[\delta]})$ and $(\mathbf{A}^{(k,\delta)}, \mathbf{B})$. On unpruned nodes, both protocols act the same. On low-value nodes, the protocols might have different control schemes, but their outputs share the same distribution. On high-value nodes, the value of $((\mathbf{A}^{[\delta]})^{(k)}, \mathbf{B}^{[\delta]})$ might be as high as 1, since $(\mathbf{A}^{[\delta]})^{(k)}$ attacks such nodes; in $(\mathbf{A}^{(k,\delta)}, \mathbf{B})$, when a high-value node is reached, $\mathbf{A}^{(k,\delta)}$ acts honestly, but since this is a high-value node, its value is at least $1-\delta$. All in all, the values of the two protocols differ by at most δ .

In the rest of this section we extend the above intuition for approximated attackers attacking approximated pruned protocols. Specifically, we give an approximated variant of $\mathbf{A}^{(k,\delta)}$ — which we call the Pruning-in-the-Head attacker — and show that it biases any protocol Π almost as well as the recursive approximated biased-continuation attacker biases the δ -approximated pruned variant of Π .

Algorithm 4.33 $(\widehat{\mathbf{A}}_{\Pi}^{(i,\xi,\delta)})$.

Input: transcript $u \in \{0, 1\}^*$.

Notation: let $\tilde{\Pi} = \Pi^{[2\delta, \xi]}$ and let $\mathcal{F} = \text{frnt} \left(\text{Small}_{\Pi}^{2\delta, \text{Est}_{\Pi}^{\xi}} \cup \text{LARGE}_{\Pi}^{2\delta, \text{Est}_{\Pi}^{\xi}} \right)$.

Operation:

1. If $u \in \mathcal{L}(\Pi)$, output $\chi_{\Pi}(u)$ and halt.
2. Set msg as follows.
 - If $u \in \text{desc}(\mathcal{F})$, set $\text{msg} = \text{HonCont}_{\Pi}^{\xi}(u)$.
 - Otherwise, set $\text{msg} = A_{\Pi}^{(i, \xi, \delta)}(u)$ (see Algorithm 4.2).
3. Send msg to B.
4. If $u' = u \circ \text{msg} \in \mathcal{L}(\Pi)$, output $\chi_{\Pi}(u')$.

The next lemma gives a lower bound on the success probability of the pruning-in-the-head attacker. It states that if a given protocol Π does not have many nodes whose value is close to 2δ , then the above algorithm (i.e., the pruning-in-the-head attacker) biases Π almost as well as the approximated attacker biases the approximated pruned protocol.

Lemma 4.34. *Let $0 < \delta \leq \delta' \leq \frac{1}{4}$, let $c = c(\delta)$ be according to Lemma 4.6, let $\xi \in (0, 1)$ and let $\tilde{\Pi} = (\tilde{A}, \tilde{B}) = \Pi^{[2\delta', \xi]}$ be the $(2\delta', \xi)$ -approximately pruned variant of an m -round protocol Π . Then*

$$\begin{aligned} \text{val} \left(\hat{A}_{\Pi}^{(k, \xi, \delta')}, \tilde{B} \right) &\geq \text{val} \left(A_{\Pi}^{(k)}, \tilde{B} \right) - 2\delta' - 2 \cdot (m+1) \cdot \sqrt{\xi} \\ &\quad - 2 \cdot \phi^{\text{Bal}}(\sqrt{\xi} + 2 \cdot m \cdot \xi, \text{border}_{\Pi}(2\delta', \xi) + 4 \cdot m \cdot \xi / \delta', \delta, \gamma) \\ &\quad - 3 \cdot \phi^{\text{Lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma), \end{aligned}$$

for every $k \in \mathbb{N}$ and $\gamma = (\gamma_1, \dots, \gamma_k)$ with $\gamma_i > 1$ for every $i \in [k]$.

4.3.1 Proving Lemma 4.34

The idea of the proof is to establish the above intuition for the pruning-in-the-head attacker, which uses an honest continuator and an estimator. This intuition indeed holds when the pruning-in-the-head attacker does not generate transcripts on which its approximating oracles fail; thus we need to bound the probability of hitting such transcripts. The probability of the original protocol to hit such failing transcripts is small, and thus we can use the results of the previous section (Proposition 4.22) to argue that this probability remains small even for the recursive approximated biased-continuation attacker, when attacking the approximated pruned protocol. Consider the following cases:

1. If the failing transcript precedes the pruned transcript (high- or low-value transcripts), then the pruning-in-the-head attacker behaves just like the recursive approximated biased-continuation attacker, so the probability remains small.
2. If the pruned transcript precedes the failing one, then we must consider two additional cases.

- (a) In the first case, the probability of the original protocol to hit a pruned transcripts is small, and in this case we are done by the above argument since, until reaching a pruned transcript, both attackers behave in exactly the same way.
- (b) In the second case, the probability of the original protocol to hit pruned transcript is high. In this case, however, the probability of the original protocol to generate failing transcripts given that the protocol reached a pruned transcript is small. Since the pruning-in-the-head attacker behaves just like the original protocol once it reaches a pruned transcript, the probability of generating failing transcripts in this case remains small.

All in all, we get that the probability that the pruning-in-the-head attacker will hit transcripts on which its approximating oracles fail is small, and thus the intuition from the ideal case applies.

Moving to the formal proof, fix δ' , ξ , k and γ as described in the statement of the lemma, and let $\tilde{\Pi} = (\tilde{\mathbf{A}}, \tilde{\mathbf{B}}) = \Pi^{[2\delta', \xi]}$. Consider the (hybrid) protocols Π_0, \dots, Π_5 . Let $\Pi_0 = (\hat{\mathbf{A}}_{\tilde{\Pi}}^{(k, \xi, \delta')}, \mathbf{B})$, $\Pi_4 = (\mathbf{A}_{\tilde{\Pi}}^{(k, \delta', \xi)}, \tilde{\mathbf{B}})$, and $\Pi_5 = (\mathbf{A}_{\tilde{\Pi}}^{(k)}, \tilde{\mathbf{B}})$. The remaining protocols are defined using the following sets:

- $\mathcal{F}ailCont^\xi := \text{frnt}(\{u \in \mathcal{V}(\Pi) : \text{SD}(\text{HonCont}\Pi^\xi(u), \text{HonCont}(u)) > \xi\})$.
- $\mathcal{L}arge := \text{frnt}(\mathcal{L}arge_{\tilde{\Pi}}^{2\delta', \xi} \setminus \text{desc}(\mathcal{F}ailCont^\xi \cup \mathcal{S}mall_{\tilde{\Pi}}^{2\delta', \xi}))$.
- $\mathcal{S}mall := \text{frnt}(\mathcal{S}mall_{\tilde{\Pi}}^{2\delta', \xi} \setminus \text{desc}(\mathcal{F}ailCont^\xi \cup \mathcal{L}arge_{\tilde{\Pi}}^{2\delta', \xi}))$.
- $\mathcal{F}ailContMid := \mathcal{F}ailCont^\xi \setminus \text{desc}(\mathcal{L}arge \cup \mathcal{S}mall)$.
- $\mathcal{F}ailContLarge := \mathcal{F}ailCont^\xi \cap \text{desc}(\mathcal{L}arge)$.
- $\mathcal{F}ailContSmall := \mathcal{F}ailCont^\xi \cap \text{desc}(\mathcal{S}mall)$.
- $\mathcal{F}ailEst := \{u \in \mathcal{V}(\Pi) : \text{val}(\Pi_u) < 1 - 2\delta' - \xi \wedge \text{Est}_{\tilde{\Pi}}^\xi(u) > 1 - 2\delta'\}$.
- $\mathcal{F}ailEstLarge := \mathcal{F}ailEst \cap \mathcal{L}arge$.
- $\mathcal{F}ailEstSmall := \mathcal{F}ailEst \cap \mathcal{S}mall$.

Note that $\mathcal{F}ailContMid$, $\mathcal{L}arge$ and $\mathcal{S}mall$ are disjoint, and that $\mathcal{F}ailContLarge$ and $\mathcal{F}ailContSmall$ are proper descendants of $\mathcal{L}arge$ and $\mathcal{S}mall$ respectively. In all the protocols described below, the control scheme and the coloring function are identical to those of Π_0 .

- Π_1 : Both parties act as in Π_4 , with the following exception the first time the parties reach a node $u \in \mathcal{F}ailContMid \cup \mathcal{L}arge \cup \mathcal{S}mall$. If $u \in \mathcal{F}ailContMid$, the parties act as in Π_0 until reaching a leaf. If $u \in \mathcal{L}arge$ the parties act as in Π , except when first reaching a node in $\mathcal{F}ailContLarge$, where the parties then act as in Π_0 until reaching a leaf. If $u \in \mathcal{S}mall$ the parties act as in Π , except when first reaching a node in $\mathcal{F}ailContSmall$, where the parties then act as in Π_0 until reaching a leaf.

- Π_2 : Both parties act as in Π_1 , with the following exception the first time the parties reach a node $u \in \mathcal{F}ailContMid \cup \mathcal{F}ailContLarge \cup \mathcal{F}ailContSmall$. If $u \in \mathcal{F}ailContMid \cup \mathcal{F}ailContSmall$, the parties act as in Π_4 until reaching a leaf. If $u \in \mathcal{F}ailContLarge$, the parties act as in Π until reaching a leaf.
- Π_3 : Both parties act as in Π_2 , with the following exception the first time the parties reach a node $u \in \mathcal{F}ailEstLarge \cup \mathcal{F}ailEstSmall$, where the parties act as in Π_4 .

The following sequence of claims, bounding the statistical distance between each pair of “neighboring” protocols, yields that $(\widehat{A}_{\Pi}^{(k,\xi,\delta')}, \mathbf{B})$ and $(A_{\Pi}^{(k)}, \widetilde{\mathbf{B}})$ are close, and the proof of the lemma follows.

Claim 4.35. *It holds that $SD(\langle \Pi_0 \rangle, \langle \Pi_1 \rangle) \leq m \cdot \xi$.*

Proof. The difference between the protocols is as follows.

- For nodes not in $\text{desc}(\mathcal{F}ailContMid \cup \mathcal{L}arge \cup \mathcal{S}mall)$: Π_0 behaves like $(A_{\Pi}^{(k,\xi,\delta')}, \mathbf{B})$; namely if u is controlled by \mathbf{A} , then the next message in Π_0 is $A_{\Pi}^{(k,\xi,\delta')}(u)$ and if u is controlled by \mathbf{B} , then the next message in Π_0 is $\text{HonCont}(u)$. Π_1 behaves like $(A_{\Pi}^{(k,\xi,\delta')}, \widetilde{\mathbf{B}})$; namely it behaves the same as Π_0 if u is controlled by \mathbf{A} , and if u is controlled by \mathbf{B} , then the next message in Π_0 is $\text{HonCont}_{\Pi}^{\xi}(u)$.
- For nodes in $\text{desc}(\mathcal{F}ailContMid)$: both protocols behave like Π_0 .
- For nodes in $\text{desc}(\mathcal{L}arge) \setminus \text{desc}(\mathcal{F}ailContLarge)$ or in $\text{desc}(\mathcal{S}mall) \setminus \text{desc}(\mathcal{F}ailContSmall)$: Π_0 acts as $(\widetilde{\mathbf{A}}, \mathbf{B})$; namely if u is controlled by \mathbf{A} , then the next message in Π_0 is $\text{HonCont}_{\Pi}^{\xi}(u)$ and if u is controlled by \mathbf{B} , then the next message in Π_0 is $\text{HonCont}_{\Pi}(u)$. Π_1 behaves like (\mathbf{A}, \mathbf{B}) ; namely the next message in Π_0 is always $\text{HonCont}_{\Pi}(u)$.
- For nodes in $\text{desc}(\mathcal{F}ailContLarge)$ or in $\text{desc}(\mathcal{F}ailContSmall)$: both protocols behave like Π_0 .

From the above case analysis, we can see that the differences between the two protocols are in nodes where one protocols calls $\text{HonCont}_{\Pi}(u)$ and the other calls $\text{HonCont}_{\Pi}^{\xi}(u)$, all for nodes u not in $\mathcal{F}ailCont^{\xi}$. Since there are at most m such calls, the claim follows. \square

Claim 4.36. *It holds that*

$$\begin{aligned} SD(\langle \Pi_1 \rangle, \langle \Pi_2 \rangle) &\leq \sqrt{\xi} + 2 \cdot \phi^{\text{Bal}}(\sqrt{\xi} + 2 \cdot m \cdot \xi, \text{border}_{\Pi}(2\delta', \xi) + 4 \cdot m \cdot \xi/\delta', \delta, \gamma) \\ &\quad + 2 \cdot \phi^{\text{lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma). \end{aligned}$$

Proof. Since Π_1 and Π_2 are identical until reaching nodes in $\mathcal{F}ailContMid \cup \mathcal{F}ailContLarge \cup \mathcal{F}ailContSmall$, it holds that

$$\begin{aligned} SD(\langle \Pi_1 \rangle, \langle \Pi_2 \rangle) &\leq \Pr_{\langle \Pi_1 \rangle} [\text{desc}(\mathcal{F}ailContMid \cup \mathcal{F}ailContLarge \cup \mathcal{F}ailContSmall)] \tag{71} \\ &\leq \Pr_{\langle \Pi_1 \rangle} [\text{desc}(\mathcal{F}ailContMid)] + \Pr_{\langle \Pi_1 \rangle} [\text{desc}(\mathcal{F}ailContLarge \cup \mathcal{F}ailContSmall)]. \end{aligned}$$

To conclude the proof we upper-bound the above two right-hand side terms. The definition of $\text{HonCont}_{\Pi}^{\xi}$ yields that both $\Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}ailContMid)]$ and

$\Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}\text{ailContLarge} \cup \mathcal{F}\text{ailContSmall})]$ are at most ξ , and the definition of Π_1 yields that

$$\begin{aligned} & \Pr_{\langle \Pi_1 \rangle} [\text{desc}(\mathcal{F}\text{ailContMid})] \\ &= \Pr_{\langle A_{\Pi}^{(i, \xi, \delta')}, \tilde{B} \rangle} [\text{desc}(\mathcal{F}\text{ailContMid})] \\ &\leq \phi^{\text{Bal}}(\xi + 2 \cdot m \cdot \xi, \text{border}_{\Pi}(2\delta', \xi) + 4 \cdot m \cdot \xi/\delta', \delta, \gamma) + \phi^{\text{lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma), \end{aligned} \quad (72)$$

where the inequality follows from Proposition 4.22. This bounds one right-hand side term of Equation (71). To bound the second term we show that

$$\begin{aligned} & \Pr_{\langle \Pi_1 \rangle} [\text{desc}(\mathcal{F}\text{ailContLarge} \cup \mathcal{F}\text{ailContSmall})] \\ &\leq \sqrt{\xi} + \phi^{\text{Bal}}(\sqrt{\xi} + 2 \cdot m \cdot \xi, \text{border}_{\Pi}(2\delta', \xi) + 4 \cdot m \cdot \xi/\delta', \delta, \gamma) + \phi^{\text{lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma), \end{aligned} \quad (73)$$

which completes the proof of the claim.

Let $\alpha = \Pr_{\langle A_{\Pi}^{(k, \xi, \delta')}, \tilde{B} \rangle} [\text{desc}(\mathcal{L}\text{arge} \cup \mathcal{S}\text{mall})]$ and $\beta = \Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}\text{ailContLarge} \cup \mathcal{F}\text{ailContSmall}) \mid \text{desc}(\mathcal{L}\text{arge} \cup \mathcal{S}\text{mall})]$. The definition of Π_1 yields that

$$\Pr_{\langle \Pi_1 \rangle} [\text{desc}(\mathcal{F}\text{ailContLarge} \cup \mathcal{F}\text{ailContSmall})] = \alpha \cdot \beta. \quad (74)$$

Assuming without loss of generality that $\beta > \sqrt{\xi}$ (as otherwise Equation (73) holds trivially) and since $\xi \geq \Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}\text{ailContLarge} \cup \mathcal{F}\text{ailContSmall})] = \Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{L}\text{arge} \cup \mathcal{S}\text{mall})] \cdot \beta$, it follows that $\Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{L}\text{arge} \cup \mathcal{S}\text{mall})] \leq \sqrt{\xi}$. Using Proposition 4.22 again, we get that

$$\alpha \leq \phi^{\text{Bal}}(\sqrt{\xi} + 2 \cdot m \cdot \xi, \text{border}_{\Pi}(2\delta', \xi) + 4 \cdot m \cdot \xi/\delta', \delta, \gamma) + \phi^{\text{lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma),$$

and Equation (73) follows since $\alpha, \beta \leq 1$ and thus $\alpha \cdot \beta \leq \alpha + \beta$. \square

Claim 4.37. *It holds that*

$$\begin{aligned} \text{SD}(\langle \Pi_2 \rangle, \langle \Pi_3 \rangle) &\leq \phi^{\text{Bal}}(\xi + 2 \cdot m \cdot \xi, \text{border}_{\Pi}(2\delta', \xi) + 4 \cdot m \cdot \xi/\delta', \delta, \gamma) \\ &\quad + \phi^{\text{lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma). \end{aligned}$$

Proof. Π_2 and Π_3 are identical until reaching nodes in $\mathcal{F}\text{ailEstLarge} \cup \mathcal{F}\text{ailEstSmall}$ and until then they both behave like $(A_{\Pi}^{(k, \xi, \delta')}, \tilde{B})$. Thus, it holds that

$$\text{SD}(\langle \Pi_2 \rangle, \langle \Pi_3 \rangle) \leq \Pr_{\langle A_{\Pi}^{(k, \xi, \delta')}, \tilde{B} \rangle} [\text{desc}(\mathcal{F}\text{ailEstLarge} \cup \mathcal{F}\text{ailEstSmall})].$$

By the definition of Est_{Π}^{ξ} , it holds that $\Pr_{\langle \Pi \rangle} [\text{desc}(\mathcal{F}\text{ailEstLarge} \cup \mathcal{F}\text{ailEstSmall})] \leq \xi$, and the claim follows Proposition 4.22. \square

Claim 4.38. *It holds that $\text{val}(\Pi_4) - \text{val}(\Pi_3) \leq 2 \cdot \delta' + (m + 1) \cdot \xi$.*

Proof. Π_3 and Π_4 are identical until reaching nodes in $\mathcal{L}\text{arge} \setminus \mathcal{F}\text{ailEstLarge}$ or $\mathcal{S}\text{mall} \setminus \mathcal{F}\text{ailEstSmall}$. Thus, it holds that

$$\text{val}(\Pi_4) - \text{val}(\Pi_3) \leq \max_{u \in \mathcal{L}\text{arge} \setminus \mathcal{F}\text{ailEstLarge} \cup \mathcal{S}\text{mall} \setminus \mathcal{F}\text{ailEstSmall}} \{\text{val}((\Pi_4)_u) - \text{val}((\Pi_3)_u)\}. \quad (75)$$

Let u be a node in which the above maximum reaches its value. The proof splits according to the set containing u .

$u \in \mathcal{L}arge \setminus \mathcal{F}ailEstLarge$: Since $u \in \mathcal{L}arge$, it holds that $\text{Est}_{\Pi}^{\xi}(u) \geq 1 - 2\delta'$, and since $u \notin \mathcal{F}ailEstLarge$, it holds that $\text{val}(\Pi_u) \geq 1 - 2\delta' - \xi$. Furthermore, the definition of Π_3 yields that $\text{val}((\Pi_3)_u) = \text{val}(\Pi_u)$, and it always holds that $\text{val}((\Pi_4)_u) \leq 1$. Thus $\text{val}((\Pi_4)_u) - \text{val}((\Pi_3)_u) \leq 2\delta' + \xi$.

$u \in \mathcal{S}mall \setminus \mathcal{F}ailEstSmall$: The protocols $(\Pi_3)_u$ and $(\Pi_4)_u$ differ only in nodes in $\text{desc}(u) \setminus \text{desc}(\mathcal{F}ailContSmall)$. If v is such a node then $(\Pi_3)_u$ behaves like Π ; namely the next message in $(\Pi_3)_u$ is $\text{HonCont}_{\Pi}(v)$. $(\Pi_4)_u$ behaves like $(A_{\Pi}^{(k,\xi,\delta')}, \tilde{B})$; namely the next message in $(\Pi_4)_u$ is $\text{HonCont}_{\Pi}^{\xi}(v)$ (u is in $\mathcal{S}mall_{\Pi}^{2\delta', \text{Est}_{\Pi}^{\xi}}$, so v is under \tilde{B} 's control, and it simply calls $\text{HonCont}_{\Pi}^{\xi}(v)$).

Since all the above calls are not in $\mathcal{F}ailContSmall$, and there are at most m such calls, it holds that $\text{SD}(\langle(\Pi_3)_u\rangle, \langle(\Pi_4)_u\rangle) \leq m \cdot \xi$. It follows that $\text{val}((\Pi_4)_u) - \text{val}((\Pi_3)_u) \leq m \cdot \xi$. □

Claim 4.39. $\text{SD}(\langle\Pi_4\rangle, \langle\Pi_5\rangle) \leq \phi^{\text{lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma)$.

Proof. This is exactly the statement of Lemma 4.21. □

4.3.2 Implementing the Pruning-in-the-Head Attacker Using an Honest Continuator

The pruning-in-the-head attacker (Algorithm 4.33) requires three algorithms: honest-continuator, estimator (both defined with respect to the attacked protocol), and the approximated biased-continuation attacker. We have already seen (Lemma 4.23) that the approximated biased-continuation attacker can be implemented using only an honest continuator. Here we take one more step and show how to implement an estimator using an honest continuator. This will almost immediately give us the implementation of the pruning-in-the-head attacker using only an honest-continuator for the attacked protocol. Specifically, the technical details of the estimator's implementation require us to slightly tweak the pruning-in-the-head attacker, so we can implement it efficiently.

From honest continuation to estimation. Turning an honest continuator into a *randomized* estimator is straightforward: given a transcript u , sample many honest continuations from u and return the mean of the parties' common outcome bit of these continuations. The pruning-in-the-head attacker, however, requires a deterministic estimator. By using standard techniques (polynomially many repetitions to get an exponentially small error, and then union bound over all possible partial transcripts) we can fix, at random, the coins of the randomized estimator in order to obtain, with high probability, a sufficiently good deterministic estimator.

Algorithm 4.40 ($\text{Est}_{\Pi}^{(\xi, \text{HC})}$).

Parameters: $\xi \in (0, 1)$.

Oracle: algorithm HC.

Input: transcript $u \in \mathcal{V}(\Pi)$.

Operation:

1. Set $\text{sum} = 0$ and $s = \left\lceil \frac{\ln\left(\frac{2^{m+1}}{\xi}\right)}{\xi^2/2} \right\rceil$.

2. For $i = 1$ to s : $sum = sum + \chi_{\Pi}(\text{HC}(u))$.
3. Return sum/s .

As before, we omit the subscript Π from the above notation of the algorithm. We also use the following convention.

Notation 4.41. Let ρ_{Est} be an upper bound on the number of random bits used by **Est** in a single call (i.e., the number of random coins used to call **HC** times s). For $r \in \{0, 1\}^{\rho_{\text{Est}}}$, let $\text{Est}_r^{(\xi, \text{HC})}$ denote the deterministic algorithm defined by hard-wiring r into the randomness of $\text{Est}^{(\xi, \delta, \text{HC})}$.

Claim 4.42. Let Π be an m -round protocol and let $\xi \in (0, 1)$. Assume that **HC** is a $\xi/2$ -honest continuator for Π . Then

$$\Pr_{r \leftarrow \{0, 1\}^{\rho_{\text{Est}}}} \left[\text{Est}_r^{(\xi, \text{HC})} \text{ is a } \xi\text{-estimator for } \Pi \right] \geq 1 - \xi.$$

Proof. For ease of notation, let $\text{Est}_r = \text{Est}_r^{(\xi, \text{HC})}$. For $u \in \mathcal{V}(\Pi)$ let $\mu_u = \mathbb{E}_{\ell \leftarrow \text{HC}(u)} [\chi(\ell)]$, and for $r \in \{0, 1\}^{\rho_{\text{Est}}}$ let A_r denote the event that $\forall u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi): |\text{Est}_r(u) - \mu_u| \leq \xi/2$. The proof is an immediate conclusion from the following two simple observations.

- (1) Est_r is a ξ -estimator for Π , $\forall r \in \{0, 1\}^{\rho_{\text{Est}}}$ for which A_r occurs.
- (2) $\Pr_{r \leftarrow \{0, 1\}^{\rho_{\text{Est}}}} [\neg A_r] \leq \xi$.

Proof of (1): Compute

$$\begin{aligned} & \Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1): |\text{Est}_r(\ell_{1, \dots, i}) - \text{val}(\Pi_{\ell_{1, \dots, i}})| > \xi] \\ & \leq \Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1): |\text{Est}_r(\ell_{1, \dots, i}) - \mu_{\ell_{1, \dots, i}}| > \xi/2 \vee |\mu_{\ell_{1, \dots, i}} - \text{val}(\Pi_{\ell_{1, \dots, i}})| > \xi/2] \\ & \leq \Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1): |\text{Est}_r(\ell_{1, \dots, i}) - \mu_{\ell_{1, \dots, i}}| > \xi/2] \\ & \quad + \Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1): |\mu_{\ell_{1, \dots, i}} - \text{val}(\Pi_{\ell_{1, \dots, i}})| > \xi/2]. \end{aligned} \tag{76}$$

Since, by assumption, A_r occurs, the first summand of the right-hand side of Equation (76) is zero. Furthermore, since **HC** is a $\xi/2$ -honest continuator for Π , we bound the second summand of the right-hand side of Equation (76):

$$\begin{aligned} & \Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1): |\mu_{\ell_{1, \dots, i}} - \text{val}(\Pi_{\ell_{1, \dots, i}})| > \xi/2] \\ & \leq \Pr_{\ell \leftarrow \langle \Pi \rangle} [\exists i \in (m-1): \text{SD}(\text{HC}(\ell_{1, \dots, i}), \text{HonCont}_{\Pi}(\ell_{1, \dots, i})) > \xi/2] \\ & \leq \xi/2 \leq \xi. \end{aligned}$$

Plugging the above into Equation (76) completes the proof.

Proof of (2): We use the following fact derived from Hoeffding's bound.

Fact 4.43 (sampling). Let $t \geq \frac{\ln(\frac{2}{\gamma})}{2 \cdot \varepsilon^2}$, let $X_1, \dots, X_t \in [0, 1]$ be iid Boolean random variables, and let $\mu = \mathbb{E}[X_i]$. Then $\Pr \left[\left| \frac{1}{t} \sum_{i=1}^t X_i - \mu \right| \geq \varepsilon \right] \leq \gamma$.

Inserting $\varepsilon := \xi/2$ and $\gamma := \xi/2^m$ into Fact 4.43 yields that

$$\Pr_{r \leftarrow \{0,1\}^{\rho_{\text{Est}}}} [|\text{Est}_r(u) - \mu_u| > \xi/2] \leq \frac{\xi}{2^m} \quad (77)$$

for every $u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi)$, and a union bound yields that

$$\begin{aligned} \Pr_{r \leftarrow \{0,1\}^{\rho_{\text{Est}}}} [\neg A_r] &= \Pr_{r \leftarrow \{0,1\}^{\rho_{\text{Est}}}} [\exists u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi) : |\text{Est}_r(u) - \mu_u| > \xi/2] \\ &\leq \sum_{u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi)} \Pr_{r \leftarrow \{0,1\}^{\rho_{\text{Est}}}} [|\text{Est}_r(u) - \mu_u| > \xi/2] \\ &\leq \sum_{u \in \mathcal{V}(\Pi) \setminus \mathcal{L}(\Pi)} \frac{\xi}{2^m} = \xi. \end{aligned}$$

□

Tweaking the pruning-in-the-head attacker. The “tweaked” pruning-in-the-head attacker invokes the pruning-in-the-head attacker, while implementing the estimator using Algorithm 4.40. Since the pruning-in-the-head attacker requires the estimator to be *deterministic*, the “tweaked” attacker fixes the estimator’s randomness for the entire execution.

Definition 4.44. Let Π be an m -round protocol and let HC be an algorithm. The algorithm $\tilde{\text{A}}_{\Pi}^{(k,\xi,\delta,\text{HC})}$ operates as follows. Before the first call to it, $\tilde{\text{A}}_{\Pi}^{(k,\xi,\delta,\text{HC})}$ sets $\text{HonCont}_{\Pi}^{\xi} := \text{HC}$ and $\text{Est}_{\Pi}^{\xi} := \text{Est}_{\Pi,r}^{(\xi,\text{HonCont}_{\Pi}^{\xi})}$, where the latter is Algorithm 4.40 when its coins are fixed to r , chosen uniformly at random. Now, when $\tilde{\text{A}}_{\Pi}^{(k,\xi,\delta,\text{HC})}$ is called with transcript u , it replies with $\hat{\text{A}}_{\Pi}^{(k,\xi,\delta)}(u)$, the answer of the pruning-in-the-head attacker from Algorithm 4.33.

Lemma 4.45. Let $\Pi = (\text{A}, \text{B})$ be an m -round protocol, let $0 < \delta \leq \delta' \leq \frac{1}{4}$ and let $\xi \in (0, 1)$. Assume that HC is a $\xi/2$ -honest continuator for Π that uses ρ_{HC} random bits and runs in time T_{HC} . Then, the algorithm $\tilde{\text{A}}_{\Pi}^{(i,\xi,\delta',\text{HC})}$ has the following properties:

1. Let $\tilde{\Pi} = (\tilde{\text{A}}, \tilde{\text{B}}) = \Pi^{[2\delta', \xi]}$ be the $(2\delta', \xi)$ -approximately pruned variant of Π . It holds that

$$\begin{aligned} \text{val} \left(\tilde{\text{A}}_{\Pi}^{(k,\xi,\delta',\text{HC})}, \text{B} \right) &\geq \text{val} \left(\text{A}_{\Pi}^{(k)}, \tilde{\text{B}} \right) - 2\delta' - 2 \cdot (m+1) \cdot \sqrt{\xi} - \xi \\ &\quad - 2 \cdot \phi^{\text{Bal}}(\sqrt{\xi} + 2 \cdot m \cdot \xi, \text{border}_{\Pi}(2\delta', \xi) + 4 \cdot m \cdot \xi/\delta', \delta, \gamma) \\ &\quad - 3 \cdot \phi^{\text{lt}}(\text{border}_{\Pi}(2\delta', \xi), \xi, m, \delta, \delta', \gamma), \end{aligned} \quad (78)$$

for every $k \in \mathbb{N}$ and $\gamma = (\gamma_1, \dots, \gamma_k)$ with $\gamma_i > 1$ for every $i \in [k]$.

2. $\tilde{\text{A}}_{\Pi}^{(k,\xi,\delta',\text{HC})}$ uses at most $k \cdot m^k \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^k \cdot \rho_{\text{HC}} + \left\lceil \frac{\ln\left(\frac{2^{m+1}}{\xi}\right)}{\xi^2/2} \right\rceil \cdot \rho_{\text{HC}}$ random bits.
3. $\tilde{\text{A}}_{\Pi}^{(k,\xi,\delta',\text{HC})}$'s running time is at most $2 \cdot k \cdot m^k \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^k \cdot \left\lceil \frac{\ln\left(\frac{2^{m+1}}{\xi}\right)}{\xi^2/2} \right\rceil \cdot T_{\text{HC}}$.

Proof. Let $\tilde{\Pi} = \Pi^{[2\delta', \xi]}$. We prove each item separately.

Proof of (1): This immediately follows from Lemma 4.34 and Claim 4.42. (Note the extra ξ term in the right-hand side of Equation (78), which accounts for the probability that $\text{Est}_{\tilde{\Pi}}^{\xi}$ is not a ξ -estimator.)

Proof of (2): The proof follows from Lemma 4.23(1). $\tilde{\mathbf{A}}_{\tilde{\Pi}}^{(k, \xi, \delta', \text{HC})}$ chooses random bits for the estimator, and then either chooses random bits for $\mathbf{A}_{\tilde{\Pi}}^{(k, \xi, \delta)}$ or $\text{HonCont}_{\tilde{\Pi}}^{\xi}$. We focus on the former case, as $\mathbf{A}_{\tilde{\Pi}}^{(k, \xi, \delta)}$ uses more random bits than $\text{HonCont}_{\tilde{\Pi}}^{\xi}$. Since $\text{HonCont}_{\tilde{\Pi}}^{\xi}$ was set to HC, the number of random bits it uses is ρ_{HC} . Thus, by Lemma 4.23(1), the number of random bits used for all calls to $\mathbf{A}_{\tilde{\Pi}}^{(k, \xi, \delta)}$ is at most $k \cdot m^k \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta'))} \right\rceil^k \cdot \rho_{\text{HC}} \leq k \cdot m^k \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^k \cdot \rho_{\text{HC}}$. Adding the number of random bits the estimator uses, which, by Algorithm 4.40, is $\left\lceil \frac{\ln\left(\frac{2^{m+1}}{\xi}\right)}{\xi^2/2} \right\rceil \cdot \rho_{\text{HC}}$, completes the proof.

Proof of (3): The proof follows from Lemma 4.23(2). $\tilde{\mathbf{A}}_{\tilde{\Pi}}^{(k, \xi, \delta', \text{HC})}$ makes a single call to $\text{Est}_{\tilde{\Pi}}^{\xi}$, and then either calls $\mathbf{A}_{\tilde{\Pi}}^{(k, \xi, \delta)}$ or $\text{HonCont}_{\tilde{\Pi}}^{\xi}$.⁵⁰ We focus on the former case, as the running time of $\mathbf{A}_{\tilde{\Pi}}^{(k, \xi, \delta)}$ is longer than that of $\text{HonCont}_{\tilde{\Pi}}^{\xi}$. Since $\text{HonCont}_{\tilde{\Pi}}^{\xi}$ was set to HC, the running time of $\text{HonCont}_{\tilde{\Pi}}^{\xi}$ is T_{HC} , and by Algorithm 4.40 the running time of $\text{Est}_{\tilde{\Pi}}^{\xi}$ is $\left\lceil \frac{\ln\left(\frac{2^{m+1}}{\xi}\right)}{\xi^2/2} \right\rceil \cdot T_{\text{HC}}$. By Lemma 4.23(2), the running time of $\tilde{\mathbf{A}}_{\tilde{\Pi}}^{(k, \xi, \delta', \text{HC})}$ is thus

$$k \cdot m^k \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta'))} \right\rceil^i \left(T_{\text{HC}} + \left\lceil \frac{\ln\left(\frac{2^{m+1}}{\xi}\right)}{\xi^2/2} \right\rceil \cdot T_{\text{HC}} \right) + \left\lceil \frac{\ln\left(\frac{2^{m+1}}{\xi}\right)}{\xi^2/2} \right\rceil \cdot T_{\text{HC}} \leq$$

$$2 \cdot k \cdot m^k \cdot \left\lceil \frac{\log(1/\xi)}{\log(1/(1-\delta))} \right\rceil^k \cdot \left\lceil \frac{\ln\left(\frac{2^{m+1}}{\xi}\right)}{\xi^2/2} \right\rceil \cdot T_{\text{HC}},$$

which completes the proof. \square

For Lemma 4.45 to be useful, we need the two last terms in Equation (78) to be small, and specifically $\text{border}_{\Pi}(2\delta', \xi)$ to be small. Proposition 4.19 yields that there is a choice for δ' such that $\text{border}_{\Pi}(2\delta', \xi)$ is small, and that this choice can be made from a polynomial-sized set. When using the above attack (see the next section), we will iterate over the polynomially-many different choices of δ' , to find a value with respect to which the above terms are indeed small.

4.4 Main Theorem — Constructing an Efficient Attacker

We are finally ready to state and prove our main result – the existence of any constant bias (even weak) coin-flipping protocol implies the existence of one-way functions.

In the following we consider both protocols and algorithms that get a security parameter, written in unary, as input (sometimes, in addition to other input), and protocols and algorithms

⁵⁰ As written in Algorithm 4.33, $\tilde{\mathbf{A}}_{\tilde{\Pi}}^{(k, \xi, \delta', \text{HC})}$ might make m calls to $\text{Est}_{\tilde{\Pi}}^{\xi}$ (checking whether $u \in \text{desc}(\mathcal{F})$ in step 2 of the algorithm). This, however, can be easily avoided by having the attacker keep a state.

that do not get a security parameter, as we did in previous sections. We refer to the former type as parametrized and to the latter type as non-parametrized. It will be clear from the context whether we consider a parametrized or non-parametrized entity. In particular, a poly-time entity whose running time is measured as a function of its security parameter is by definition parametrized. Given a parametrized protocol Π and $n \in \mathbb{N}$, let Π_n be its non-parametrized variant with the security parameter 1^n hardwired into the parties' code. We apply similar notation also for parametrized algorithms.

Theorem 4.46 (main theorem, restatement of Theorem 1.1). *Assume one-way functions do not exist. Then for every PPT coin-flipping protocol $\Pi = (\mathbf{A}, \mathbf{B})$ and $\varepsilon > 0$, there exist PPTM's \mathcal{A} and \mathcal{B} such that the following hold for infinitely many n 's.*

1. $\Pr [\text{out}(\mathcal{A}(1), \mathbf{B})(1^n)] \geq 1 - \varepsilon$ or $\Pr [\text{out}(\mathbf{A}, \mathcal{B}(0))(1^n)] \leq \varepsilon$, and
2. $\Pr [\text{out}(\mathcal{A}(0), \mathbf{B})(1^n)] \leq \varepsilon$ or $\Pr [\text{out}(\mathbf{A}, \mathcal{B}(1))(1^n)] \geq 1 - \varepsilon$.

The proof of Theorem 4.46 follows from Theorem 3.3 and Lemma 4.45 together with the following lemma, which shows how to implement an efficient honest continuator assuming OWFs do not exist.

Lemma 4.47. *Assume one-way functions do not exist. Then for any PPT coin-flipping protocol $\Pi = (\mathbf{A}, \mathbf{B})$ and $p \in \text{poly}$, there exists a PPTM algorithm HC such that HC_n is a $1/p(n)$ -honest continuator for Π_n for infinitely many n 's.*

The proof of Lemma 4.47 is given below, but we first we use it to prove Theorem 4.46.

Proof of Theorem 4.46. We focus on proving the first part of the theorem, and the second, symmetric part follows the same arguments.

Let $\delta = \varepsilon/8$, and let $\xi(n) = 1/p(n) < \frac{(2\delta)^2}{16m(n)^2}$ for some large enough $p \in \text{poly}$ to be determined by the analysis. Let HC be the algorithm guaranteed by Lemma 4.47, such that HC_n is an $\xi(n)/2$ -honest continuator for Π_n for every n in an infinite set $\mathcal{I} \subseteq \mathbb{N}$. For $n \in \mathcal{I}$, let $\delta'_n \in [\delta/2, \delta]$ be such that $\text{border}_{\Pi_n}(2\delta'_n, \xi(n)) \leq m(n) \cdot \sqrt{2\xi(n)}$, guaranteed to exist from Proposition 4.19.⁵¹ Let $\tilde{\Pi}_n = (\tilde{\mathbf{A}}_n, \tilde{\mathbf{B}}_n) = \Pi_n^{[2\delta'_n, \xi]}$ be the $(2\delta'_n, \xi)$ -approximately pruned variant of Π_n and let $\kappa = \kappa(\varepsilon/2)$ be such that $\text{val}(\mathbf{A}_{\tilde{\Pi}_n}^{(k)}, \tilde{\mathbf{B}}_n) > 1 - \varepsilon/2$ or $\text{val}(\tilde{\mathbf{A}}_n, \mathbf{B}_{\tilde{\Pi}_n}^{(k)}) < \varepsilon/2$, guaranteed to exist for every $n \in \mathcal{I}$ from Theorem 3.3. Assume without loss of generality that there exists an infinite set $\mathcal{I}' \subseteq \mathcal{I}$ such that

$$\text{val}(\mathbf{A}_{\tilde{\Pi}_n}^{(k)}, \tilde{\mathbf{B}}_n) > 1 - \varepsilon/2 \tag{79}$$

for every $n \in \mathcal{I}'$, and let $c = c(\delta/2)$ from Lemma 4.6. Note that the bound in Lemma 4.6 holds for any $\delta' \geq \delta/2$ as well. Let $\gamma = (\gamma_1, \dots, \gamma_\kappa)$ be such that $\gamma_i \in \text{poly}$ for every i , to be determined by the analysis, and let $\gamma_n = (\gamma_1(n), \dots, \gamma_\kappa(n))$. We recall that $\kappa \in \mathbb{N}$ is *constant* depending only on ε from Theorem 3.3, and not a function of n .

⁵¹By the choice of ξ and by Proposition 4.19 there exists $\delta'' \in [\delta, 2\delta]$ such that $\text{border}_{\Pi_n}(\delta'', \xi(n)) \leq m(n) \cdot \sqrt{2\xi(n)}$. Now we can set $\delta' = \delta''/2$.

After preparing the background, we are now ready to determine the rest of the parameters. Compute

$$\begin{aligned} & \phi^{\text{lt}}(\text{border}_{\Pi_n}(2\delta'_n, \xi(n)), \xi(n), m(n), \delta, \delta'_n, \gamma_n) \\ &= \kappa \cdot \frac{30^\kappa \cdot m(n)^\kappa \cdot \prod_{i=1}^\kappa \gamma_i(n)}{\delta'_n{}^{2\kappa}} \cdot \left(m(n) \cdot \sqrt{2\xi(n)} + \frac{9 \cdot m(n) \cdot \xi(n)}{\delta'_n} \right) \end{aligned} \quad (80)$$

$$+ \sum_{i=1}^\kappa 2^{\kappa-i+2} \cdot \frac{30^{\kappa-i} \cdot m(n)^{\kappa-i} \cdot \prod_{j=i+1}^\kappa \gamma_j(n)}{\delta'_n{}^{2(\kappa-i)} \cdot \gamma_i(n)^c}. \quad (81)$$

Set $\gamma_\kappa \in \text{poly}$ such that $4/\gamma_\kappa(n)^c \in o(1)$ (all asymptotic notations are with respect to n). For $i \in [\kappa - 1]$, set $\gamma_i \in \text{poly}$ such that $2^{\kappa-i+2} \cdot \frac{30^{\kappa-i} \cdot m(n)^{\kappa-i} \cdot \prod_{j=i+1}^\kappa \gamma_j(n)}{\delta'_n{}^{2(\kappa-i)} \cdot \gamma_i(n)^c} \in o(1)$ (this can be done by first setting $\gamma_{\kappa-1}$, then $\gamma_{\kappa-2}$, and so on). Since κ and c are fixed and independent of n , it is guaranteed that such settings for γ exist, and that the term in Equation (81) is in $o(1)$. After setting γ , and since δ'_n is bounded in $[\delta/2, \delta]$ and κ and c are independent of n , the term in Equation (80) can be bound by $\sqrt{\xi(n)} \cdot \text{poly}(n)$. Hence, we can now set $\xi(n) = 1/p(n)$, such that $p \in \text{poly}$ and the term in Equation (80) is also in $o(1)$. Compute

$$\begin{aligned} & \phi^{\text{Bal}}(\sqrt{\xi(n)} + 2 \cdot m(n) \cdot \xi(n), \text{border}_{\Pi_n}(2\delta'_n, \xi(n)) + 4 \cdot m(n) \cdot \xi(n)/\delta'_n, \delta/2, \gamma_n) \\ &= \left(\sqrt{\xi(n)} + 2 \cdot m(n) \cdot \xi(n) + m(n) \cdot \sqrt{2\xi(n)} + \frac{4 \cdot m(n) \cdot \xi(n)}{\delta'_n} \right) \cdot \prod_{i=1}^\kappa \gamma_i(n) + 2 \cdot \sum_{i=1}^\kappa \frac{\prod_{j=i+1}^\kappa \gamma_j(n)}{\gamma_i(n)^c}. \end{aligned} \quad (82)$$

By our choice of parameters the right-hand side (and thus the left-hand side) of Equation (82) is in $o(1)$. Hence, Lemma 4.45(1) yields that

$$\text{val} \left(\tilde{\mathbf{A}}_{\Pi_n}^{(\kappa, \xi(n), \delta'_n, \text{HC}_n)}, \mathbf{B}_{\Pi_n} \right) \geq \text{val} \left(\mathbf{A}_{\Pi_n}^{(k)}, \tilde{\mathbf{B}}_n \right) - 3\delta' - o(1) \geq 1 - \frac{\varepsilon}{2} - \frac{\varepsilon}{4} - o(1). \quad (83)$$

Our final adversary $\mathcal{A}(1)$ is defined as follows: on input 1^n , it checks all possible candidates for δ'_n from Proposition 4.19, estimates the value of $\tilde{\Pi}_{\delta'_n} := \left(\tilde{\mathbf{A}}_{\Pi_n}^{(\kappa, \xi(n), \delta'_n, \text{HC}_n)}, \mathbf{B}_{\Pi_n} \right)$ by running the latter for polynomially-many times, sets δ_n^* to be the value that maximizes $\tilde{\Pi}_{\delta'_n}$, and then, when interacting with \mathbf{B} , it behaves like $\tilde{\mathbf{A}}_{\Pi_n}^{(\kappa, \xi, \delta_n^*, \text{HC}_n)}$. Since $\mathcal{A}(1)$ estimates the value of $\tilde{\Pi}_{\delta'_n}$ for polynomial many times, it estimates the value of $\tilde{\Pi}_{\delta_n^*}$ to be at least $1 - 3\varepsilon/4 - o(1)$ with exponentially small probability. Thus,

$$\Pr [\text{out}(\mathcal{A}(1), \mathbf{B})(1^n)] \geq \text{val} \left(\hat{\mathbf{A}}_{\Pi_n}^{(\kappa, \delta'_n, \xi(n), \text{HC}_n)}, \mathbf{B}_{\Pi_n} \right) - o(1) \geq 1 - 3\varepsilon/4 - o(1) \geq 1 - \varepsilon \quad (84)$$

for large enough $n \in \mathcal{I}'$.

The last step is to argue that $\mathcal{A}(1)$ is efficient. By our choice of parameters, the fact that κ is constant (i.e., independent of n) and HC is PPTM, Lemma 4.45(2,3) yields that $\tilde{\mathbf{A}}_{\Pi_n}^{(\kappa, \xi(n), \delta'_n, \text{HC}_n)}$ is a PPTM. Since there are only $\text{poly}(n)$ possibilities for setting δ'_n , it follows that the running time of $\mathcal{A}(1)$ is also is $\text{poly}(n)$. \square

It is left to prove Lemma 4.47.

Proof of Lemma 4.47. Let $m(n) = \text{round}(\Pi_n)$, and let $\rho_A(n)$ and $\rho_B(n)$ be, respectively, the (maximal) number of random bits used by A and B on common input 1^n . Consider the *transcript function* f_Π over $1^* \times \{0, 1\}^{\rho_A(n)} \times \{0, 1\}^{\rho_B(n)} \times (m(n) - 1)$, defined by

$$f_\Pi(1^n, r_A, r_B, i) = 1^n, \text{trans}((A(\cdot; r_A), B(\cdot; r_B))(1^n))_{1, \dots, i}. \quad (85)$$

Since Π is a polynomial time protocol, it follows without loss of generality that $m(n), \rho_A(n), \rho_B(n) \in \text{poly}(n)$ and that f_Π is computable in polynomial time.

Under the assumption that OWFs do not exist, the transcript function is not distributional one-way, i.e., it has an inverter that returns a random preimage. We would like to argue that an algorithm that outputs the transcript induced by the randomness this inverter returns is an honest continuator. This is almost true, as this inverter guarantees to work for a random node of the protocol tree, and we require that an honest continuator work for all nodes in a random *path* of the protocol tree. Still, since any path in the protocol tree is of polynomial length, the lemma follows by a union bound. We now move to the formal proof.

Fix $p \in \text{poly}$ and let Inv be the $1/(m \cdot p)$ -inverter guaranteed to exist by Lemma 2.12. Namely, $\text{Inv}_n = \text{Inv}(1^n, \cdot)$ is a $1/(m(n) \cdot p(n))$ -inverter for $f_\Pi(1^n, \cdot, \cdot, \cdot)$ for every n within an infinite size index set $\mathcal{I} \subseteq \mathbb{N}$.⁵² By the definition of f_Π , choosing a random preimage from $f^{-1}(1^n, u)$ is equivalent to choosing an element according to the distribution $(\text{Cosis}_{\Pi_n}(u), |u|)$.⁵³ For a transcript u , let $f_u(x, y, z) := u \circ (\text{trans}(A(\cdot; x), B(\cdot; y))(1^n))_{|u|+1, \dots, m(n)}$, and let HC_n be the algorithm that, given input u , returns $f_u(\text{Inv}_n(u))$. We show that HC_n is a $1/p(n)$ -honest continuator for Π_n , for every $n \in \mathcal{I}$.

Fix $n \in \mathcal{I}$. Let $m = m(n)$, $p = p(n)$ and from now on we omit n from notations. Note that $f_u(\text{Cosis}_\Pi(u), |u|) \equiv \langle \Pi_u \rangle \equiv \text{HonCont}_\Pi(u)$, and thus

$$\text{SD}(\text{Inv}(u), (\text{Cosis}_\Pi(u), |u|)) \geq \text{SD}(\text{HC}(u), \text{HonCont}_\Pi(u)), \quad (86)$$

for every transcript u . Let I and L be random variables distributed as $I \leftarrow (m - 1)$ and $L \leftarrow \langle \Pi \rangle$

⁵²Lemma 2.12 is stated for functions whose domain is $\{0, 1\}^n$ for every $n \in \mathbb{N}$, i.e., functions defined for every input length. Although the transcript function is not defined for every input length (and has 1^n as an input), using the fact that it is defined on $\{0, 1\}^{q(n)}$ for some $q(n) \in \text{poly}(n)$ and standard padding techniques, Lemma 2.12 does in fact guarantee such an inverter.

⁵³Recall that $\text{Cosis}_\Pi(u)$ returns random coins for the parties, consistent with a random execution of Π leading to u .

respectively. Compute

$$\begin{aligned}
& \Pr \left[\text{SD}(\text{Inv}(L_{1,\dots,I}), (\text{Consis}_{\Pi}(L_{1,\dots,I}), I)) > \frac{1}{m \cdot p} \right] \\
&= \sum_{j=0}^{m-1} \Pr \left[\text{SD}(\text{Inv}(L_{1,\dots,I}), (\text{Consis}_{\Pi}(L_{1,\dots,I}), I)) > \frac{1}{m \cdot p} \mid I = j \right] \cdot \Pr[I = j] \\
&= \frac{1}{m} \sum_{j=0}^{m-1} \Pr \left[\text{SD}(\text{Inv}(L_{1,\dots,j}), (\text{Consis}_{\Pi}(L_{1,\dots,j}), j)) > \frac{1}{m \cdot p} \right] \\
&\geq \frac{1}{m} \sum_{j=0}^{m-1} \Pr \left[\text{SD}(\text{HC}(L_{1,\dots,j}), \text{HonCont}_{\Pi}(L_{1,\dots,j})) > \frac{1}{m \cdot p} \right] \\
&\geq \frac{1}{m} \sum_{j=0}^{m-1} \Pr \left[\text{SD}(\text{HC}(L_{1,\dots,j}), \text{HonCont}_{\Pi}(L_{1,\dots,j})) > \frac{1}{p} \right] \\
&\geq \frac{1}{m} \Pr \left[\exists j \in (m-1): \text{SD}(\text{HC}(L_{1,\dots,j}), \text{HonCont}_{\Pi}(L_{1,\dots,j})) > \frac{1}{p} \right].
\end{aligned}$$

The proof now follows by the properties of `Inv`. □

References

- [1] B. Averbuch, M. Blum, B. Chor, S. Goldwasser, and S. Micali. How to implement Bracha's $O(\log n)$ Byzantine agreement algorithm, 1985. Unpublished manuscript.
- [2] A. Beimel, E. Omri, and I. Orlov. Protocols for multiparty coin toss with dishonest majority. In *Advances in Cryptology – CRYPTO 2010*, pages 538–557, 2010.
- [3] M. Blum. Coin flipping by telephone. In *Advances in Cryptology – CRYPTO '81*, pages 11–15, 1981.
- [4] A. Chailloux and I. Kerenidis. Optimal quantum strong coin flipping. In *Proceedings of the 50th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 527–533, 2009.
- [5] R. Cleve. Limits on the security of coin flips when half the processors are faulty. In *Proceedings of the 18th Annual ACM Symposium on Theory of Computing (STOC)*, pages 364–369, 1986.
- [6] R. Cleve and R. Impagliazzo. Martingales, collective coin flipping and discrete control processes (extended abstract). <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.51.1797>, 1993.
- [7] D. Dachman-Soled, Y. Lindell, M. Mahmoody, and T. Malkin. On the black-box complexity of optimally-fair coin tossing. In *Theory of Cryptography, 8th Theory of Cryptography Conference (TCC)*, volume 6597, pages 450–467, 2011.
- [8] O. Goldreich and L. A. Levin. A hard-core predicate for all one-way functions. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 25–32, 1989.

- [9] O. Goldreich, S. Goldwasser, and S. Micali. On the cryptographic applications of random functions. In *Advances in Cryptology – CRYPTO '84*, pages 276–288, 1984.
- [10] O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *Journal of the ACM*, pages 792–807, 1986.
- [11] I. Haitner and E. Omri. Coin Flipping with Constant Bias Implies One-Way Functions. In *Proceedings of the 52nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 110–119, 2011.
- [12] I. Haitner, M. Nguyen, S. J. Ong, O. Reingold, and S. Vadhan. Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM Journal on Computing*, 39(3):1153–1218, 2009.
- [13] J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, pages 1364–1396, 1999. Preliminary versions in *STOC'89* and *STOC'90*.
- [14] R. Impagliazzo. Pseudo-random generators for cryptography and for randomized algorithms. <http://cseweb.ucsd.edu/~russell/format.ps>. Ph.D. Thesis.
- [15] R. Impagliazzo and M. Luby. One-way functions are essential for complexity based cryptography. In *Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 230–235, 1989.
- [16] A. Y. Kitaev. Quantum coin-flipping. Presentation at the 6th Workshop on Quantum Information Processing (QIP 2003), 2003.
- [17] H. K. Maji, M. Prabhakaran, and A. Sahai. On the Computational Complexity of Coin Flipping. In *Proceedings of the 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 613–622, 2010.
- [18] C. Mochon. Quantum weak coin flipping with arbitrarily small bias. arXiv:0711.4114, 2007.
- [19] T. Moran, M. Naor, and G. Segev. An optimally fair coin toss. In *Theory of Cryptography, 6th Theory of Cryptography Conference (TCC)*, pages 1–18, 2009.
- [20] M. Naor. Bit commitment using pseudorandomness. *Journal of Cryptology*, pages 151–158, 1991. Preliminary version in *CRYPTO'89*.
- [21] M. Naor and M. Yung. Universal one-way hash functions and their cryptographic applications. In *Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC)*, pages 33–43, 1989.
- [22] A. W. Roberts and D. E. Varberg. *Convex Functions*. Academic Press Inc, 1973.
- [23] J. Rompel. One-way functions are necessary and sufficient for secure signatures. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 387–394, 1990.
- [24] S. Zachos. Probabilistic Quantifiers, Adversaries, and Complexity Classes: An Overview. In *Proceedings of the First Annual IEEE Conference on Computational Complexity*, pages 383–400, 1986.

A Missing Proofs

A.1 Proving Lemma 2.17

Lemma A.1 (Restatement of Lemma 2.17). *Let $x, y \in [0, 1]$ and $a_1, \dots, a_k, b_1, \dots, b_k \in (0, 1]$. Then for any $p_0, p_1 \geq 0$ with $p_0 + p_1 = 1$, it holds that*

$$p_0 \cdot \frac{x^{k+1}}{\prod_{i=1}^k a_i} + p_1 \cdot \frac{y^{k+1}}{\prod_{i=1}^k b_i} \geq \frac{(p_0 x + p_1 y)^{k+1}}{\prod_{i=1}^k (p_0 a_i + p_1 b_i)}. \quad (87)$$

Proof. The lemma easily follows if one of the following holds: (1) $p_0 = 1, p_1 = 0$; (2) $p_0 = 0, p_1 = 1$; and (3) $x = y = 0$. Assuming $1 > p_0, p_1 > 0$ and $x + y > 0$, dividing Equation (87) by its right-hand side (which is always positive) gives

$$p_0 \cdot \frac{\left(\frac{x}{p_0 x + p_1 y}\right)^{k+1}}{\prod_{i=1}^k \frac{a_i}{p_0 a_i + p_1 b_i}} + p_1 \cdot \frac{\left(\frac{y}{p_0 x + p_1 y}\right)^{k+1}}{\prod_{i=1}^k \frac{b_i}{p_0 a_i + p_1 b_i}} \geq 1. \quad (88)$$

Define the following variable changes:

$$z = \frac{p_0 x}{p_0 x + p_1 y} \quad c_i = \frac{p_0 a_i}{p_0 a_i + p_1 b_i} \quad \text{for } 1 \leq i \leq k.$$

It follows that

$$1 - z = \frac{p_1 y}{p_0 x + p_1 y} \quad 1 - c_i = \frac{p_1 b_i}{p_0 a_i + p_1 b_i} \quad \text{for } 1 \leq i \leq k.$$

Note that $0 \leq z \leq 1$ and that $0 < c_i < 1$ for every $1 \leq i \leq k$. Plugging the above into Equation (88), it remains to show that

$$\frac{z^{k+1}}{\prod_{i=1}^k c_i} + \frac{(1-z)^{k+1}}{\prod_{i=1}^k (1-c_i)} \geq 1 \quad (89)$$

for all $0 \leq z \leq 1$ and $0 < c_i < 1$. Equation (89) immediately follows for $z = 0, 1$, and in the rest of the proof we show that it also holds for $z \in (0, 1)$. Define $f(z, c_1, \dots, c_k) := \frac{z^{k+1}}{\prod_{i=1}^k c_i} + \frac{(1-z)^{k+1}}{\prod_{i=1}^k (1-c_i)} - 1$. Equation (89) follows by showing that $f(z, c_1, \dots, c_k) \geq 0$ for all $z \in (0, 1)$ and $0 < c_i < 1$. Taking the partial derivative with respect to c_i for $1 \leq i \leq k$, it holds that

$$\frac{\partial}{\partial c_i} f = -\frac{z^{k+1}}{c_i^2 \prod_{\substack{1 \leq j \leq k \\ j \neq i}} c_j} + \frac{(1-z)^{k+1}}{(1-c_i)^2 \prod_{\substack{1 \leq j \leq k \\ j \neq i}} (1-c_j)}.$$

Fix $0 < z < 1$, and let $f_z(c_1, \dots, c_k) = f(z, c_1, \dots, c_k)$. If $c_1 = \dots = c_k = z$, then for every $1 \leq i \leq k$ it holds that $\frac{\partial}{\partial c_i} f_z(c_1, \dots, c_k) = \frac{\partial}{\partial c_i} f(z, c_1, \dots, c_k) = 0$. Hence, f_z has a local extremum at $(c_1, \dots, c_k) = (z, \dots, z)$. Taking the second partial derivative with respect to c_i for $1 \leq i \leq k$, it holds that

$$\frac{\partial^2}{\partial c_i^2} f = \frac{2z^{k+1}}{c_i^3 \prod_{\substack{1 \leq j \leq k \\ j \neq i}} c_j} + \frac{2(1-z)^{k+1}}{(1-c_i)^3 \prod_{\substack{1 \leq j \leq k \\ j \neq i}} (1-c_j)} > 0,$$

and thus, $(c_1, \dots, c_k) = (z, \dots, z)$ is a local minimum of f_z .

The next step is to show that $(c_1, \dots, c_k) = (z, \dots, z)$ is a global minimum of f_z . This is done by showing that f_z is convex when $0 < c_i < 1$. Indeed, consider the function $-\ln(x)$. This is a convex function in for $0 < x < 1$. Thus the function $\sum_{i=1}^k -\ln(c_i)$, which is a sum of convex functions, is also convex. Moreover, consider the function e^x . This is a convex function for any x . Hence, the function $e^{\sum_{i=1}^k -\ln(c_i)} = \frac{1}{\prod_{i=1}^k c_i}$, which is a composition of two convex functions, is also convex for $0 < c_i < 1$. Since z is fixed, the function $\frac{z^{k+1}}{\prod_{i=1}^k c_i}$ is also convex. Similar argument shows that $\frac{(1-z)^{k+1}}{\prod_{i=1}^k (1-c_i)}$ is also convex for $0 < c_i < 1$. This yields that f_z , which is a sum of two convex functions, is convex. It is known that a local minimum of a convex function is also a global minimum for that function [22, Theorem A, Chapter V], and thus (z, \dots, z) is a global minimum of f_z .

Let $z', c'_1, \dots, c'_k \in (0, 1)$. Since (z', \dots, z') is a global minimum of $f_{z'}$, it holds that $f(z', z', \dots, z') = f_{z'}(z', \dots, z') \leq f_{z'}(c'_1, \dots, c'_k) = f(z', c'_1, \dots, c'_k)$. But $f(z', z', \dots, z') = 0$, and thus $f(z', c'_1, \dots, c'_k) \geq 0$. This shows that Equation (89) holds, and the proof is concluded. \square

A.2 Proving Lemma 2.18

Lemma A.2 (Restatement of Lemma 2.18). *For every $\delta \in (0, \frac{1}{2}]$, there exists $\alpha = \alpha(\delta) \in (0, 1]$ such that*

$$\lambda \cdot a_1^{1+\alpha} \cdot (2 - a_1 \cdot x) + a_2^{1+\alpha} \cdot (2 - a_2 \cdot x) \leq (1 + \lambda) \cdot (2 - x), \quad (90)$$

for every $x \geq \delta$ and $\lambda, y \geq 0$ with $\lambda y \leq 1$, for $a_1 = 1 + y$ and $a_2 = 1 - \lambda y$.

Proof. Fix $\delta \in (0, \frac{1}{2}]$. Rearranging the terms of Equation (90), one can equivalently prove that for some $\alpha \in (0, 1]$, it holds that

$$x \cdot (1 + \lambda - \lambda \cdot (1 + y)^{2+\alpha} - (1 - \lambda y)^{2+\alpha}) \leq 2 \cdot (1 + \lambda - \lambda \cdot (1 + y)^{1+\alpha} - (1 - \lambda y)^{1+\alpha}) \quad (91)$$

for all x, λ and y in the proper range. Note that the above trivially holds, regardless of the choice of $\alpha \in (0, 1]$, if $\lambda y = 0$ (both sides of the inequality are 0). In the following we show that for the cases $\lambda y = 1$ and $\lambda y \in (0, 1)$, Equation (91) holds for any small enough choice of α . Hence, the proof follows by taking the small enough α for which the above cases hold simultaneously.

$\lambda y = 1$: Let $z = \frac{1}{\lambda} + 1 = y + 1 > 1$. Plugging in Equation (91), we need to find $\alpha_h \in (0, 1]$ for which it holds that

$$x \cdot \left(1 + \frac{1}{z-1} - \frac{z^{2+\alpha}}{z-1}\right) \leq 2 \cdot \left(1 + \frac{1}{z-1} - \frac{z^{1+\alpha}}{z-1}\right) \quad (92)$$

for for all $z > 1$ and $\alpha \in (0, \alpha_h)$. Equivalently, by multiplying both sides by $\frac{z-1}{z}$ – which, since $z > 1$, is always positive – it suffices to find $\alpha_h \in (0, 1]$ for which it holds that

$$x \cdot (1 - z^{1+\alpha}) \leq 2 \cdot (1 - z^\alpha) \quad (93)$$

for all $z > 1$ and $\alpha \in (0, \alpha_h)$.

Since $1 - z^{1+\alpha} < 0$ for all $\alpha \geq 0$ and $z > 1$, and letting $h_\alpha(z) := \frac{z^\alpha - 1}{z^{1+\alpha} - 1}$, proving Equation (93) is equivalent to finding $\alpha_h \in (0, 1]$ such that

$$\delta \geq \sup_{z>1} \{2 \cdot h_\alpha(z)\} = 2 \cdot \sup_{z>1} \{h_\alpha(z)\} \quad (94)$$

for all $z > 1$ and $\alpha \in (0, \alpha_h)$.

Consider the function

$$h(w) := \sup_{z>1} \{h_w(z)\}. \quad (95)$$

Claim A.3 states that $\lim_{w \rightarrow 0^+} h(w) = 0$ (i.e., $h(w)$ approaches 0 when w approaches 0 from the positive side), and hence $2 \cdot \lim_{w \rightarrow 0^+} h(w) = 0$. The proof of Equation (94), and thus the proof of this part, follows since there is now small enough $\alpha_h < 1$ for which $x \geq 2 \cdot h(\alpha)$ for every $\alpha \in (0, \alpha_h]$ and $x \geq \delta$.

$\lambda y \in (0, 1)$: Consider the function

$$g(\alpha, \lambda, y) := 1 + \lambda - \lambda \cdot (1 + y)^{2+\alpha} - (1 - \lambda y)^{2+\alpha}. \quad (96)$$

Claim A.4 states that for $\alpha \geq 0$, the function g is negative over the given range of λ and y . This allows us to complete the proof by finding $\alpha \in (0, 1]$ for which

$$\delta \geq 2 \cdot \sup_{\lambda, y > 0, \lambda y < 1} \left\{ f_\alpha(\lambda, y) := \frac{1 + \lambda - \lambda \cdot (1 + y)^{1+\alpha} - (1 - \lambda y)^{1+\alpha}}{1 + \lambda - \lambda \cdot (1 + y)^{2+\alpha} - (1 - \lambda y)^{2+\alpha}} \right\}. \quad (97)$$

Consider the function

$$f(w) := \sup_{\lambda, y > 0, \lambda y < 1} \{f_w(\lambda, y)\}. \quad (98)$$

Claim A.5 states that $\lim_{w \rightarrow 0^+} h(w) = 0$, and hence $(1 + \delta) \cdot \lim_{w \rightarrow 0^+} h(w) = 0$. The proof of Equation (97), and thus the proof of this part follows since there is now small enough $\alpha_f < 1$ for which $x \geq 2 \cdot h(\alpha)$ for every $\alpha \in (0, \alpha_f]$ and $x \geq \delta$.

By setting $\alpha_{\min} = \min\{\alpha_h, \alpha_f\}$, it follows that $x \geq h(\alpha), f(\alpha)$ for any $\alpha \in (0, \alpha_{\min})$ and $x \geq \delta$, concluding the the proof of the claim. \square

Claim A.3. $\lim_{w \rightarrow 0^+} h(w) = 0$.

Proof. Simple calculations show that for fixed w , the function $h_w(z)$ is decreasing in the interval $(1, \infty)$. Indeed, fix some $w > 0$, and consider the derivative of h_w

$$\begin{aligned} h'_w(z) &= \frac{wz^{w-1}(z^{1+w} - 1) - (1+w)z^w(z^w - 1)}{(z^{1+w} - 1)^2} \\ &= \frac{-z^{w-1}(z^{1+w} - (1+w)z + w)}{(z^{1+w} - 1)^2}. \end{aligned} \quad (99)$$

Let $p(z) := z^{1+w} - (1+w)z + w$. Taking the derivative of p and equaling it to 0, we have that

$$\begin{aligned} p'(z) &= (1+w)z^w - (1+w) = 0 \\ &\iff z = 1. \end{aligned} \tag{100}$$

Since $p''(1) = (1+w)w > 0$ for all $w > 0$, it holds that $z = 1$ is the minimum of p in $[1, \infty)$. Since $p(1) = 0$, it holds that $p(a) > 0$ for every $a \in (1, \infty)$. Thus, $h'_w(z) < 0$, and $h_w(z)$ is decreasing in the interval $(1, \infty)$. The latter fact yields that

$$\begin{aligned} \lim_{w \rightarrow 0^+} h(w) &= \lim_{w \rightarrow 0^+} \sup_{z > 1} h_w(z) \\ &= \lim_{w \rightarrow 0^+} \lim_{z \rightarrow 1^+} \frac{z^w - 1}{z^{1+w} - 1} \\ &= \lim_{w \rightarrow 0^+} \lim_{z \rightarrow 1^+} \frac{wz^{w-1}}{(1+w)z^w} \\ &= \lim_{w \rightarrow 0^+} \frac{w}{1+w} \\ &= 0, \end{aligned}$$

where the third equality holds by L'Hôpital's rule. \square

Claim A.4. For all $\alpha \geq 0$ and $\lambda, y > 0$ with $\lambda y < 1$, it holds that $g(\alpha, \lambda, y) < 0$.

Proof. Fix $\lambda, y > 0$ with $\lambda y \leq 1$ and let $f(x) := g(x, \lambda, y)$. We first prove that f is strictly decreasing in the range $[0, \infty)$, and then show that $f(0) < 0$, yielding that $g(\alpha, \lambda, y) < 0$ for the given range of parameters. Taking the derivative of f , we have that

$$f'(x) = -\lambda \cdot (1+y)^{2+x} \cdot \ln(1+y) + (1-\lambda y)^{2+x} \cdot \ln(1-\lambda y), \tag{101}$$

and since $\ln(1-\lambda y) < 0$, it holds that f' is a negative function. Hence, f is strictly decreasing, and takes its (unique) maximum over $[0, \infty)$ at 0. We conclude the proof by noting that $f(0) = -\lambda \cdot y^2 \cdot (1+\lambda) < 0$. \square

Claim A.5. $\lim_{w \rightarrow 0^+} f(w) = 0$.

Proof. Assume towards a contradiction that the claim does not hold. It follows that there exist $\varepsilon > 0$ and an infinite sequence $\{w_i\}_{i \in \mathbb{N}}$ such that $\lim_{i \rightarrow \infty} w_i = 0$ and $f(w_i) \geq \varepsilon$ for every $i \in \mathbb{N}$. Hence, there exists an infinite sequence of pairs $\{(\lambda_i, y_i)\}_{i \in \mathbb{N}}$, such that for every $i \in \mathbb{N}$ it holds that $f(w_i) = f_{w_i}(\lambda_i, y_i) \geq \varepsilon$, $\lambda_i, y_i > 0$ and $\lambda_i y_i \leq 1$.

If $\{\lambda_i\}_{i \in \mathbb{N}}$ is not bounded from above, we focus on a subsequence of $\{(\lambda_i, y_i)\}$ in which λ_i converges to ∞ , and let $\lambda^* = \infty$. Similarly, if $\{y_i\}_{i \in \mathbb{N}}$ is not bounded from above, we focus on a subsequence of $\{(\lambda_i, y_i)\}$ in which y_i converges to ∞ , and let $y^* = \infty$. Otherwise, by the Bolzano-Weierstrass Theorem, there exists a subsequence of $\{(\lambda_i, y_i)\}$ in which both λ_i and y_i converge to some real values. We let λ^* and y^* be these values.

The rest of the proof splits according to the values of λ^* and y^* . In each case we focus on the subsequence of $\{(w_i, \lambda_i, y_i)\}$ that converges to $(0, \lambda^*, y^*)$, and show that $\lim_{i \rightarrow \infty} f_{w_i}(\lambda_i, y_i) = 0$, in contradiction to the above assumption.

$y^* = \infty$: First note that the assumption $y^* = \infty$ and the fact that $\lambda_i y_i \leq 1$ for every i yield that $\lambda^* = 0$.

For $c \in [0, 1)$, the Taylor expansion with Lagrange remainder over the interval $[0, c]$ yields that

$$(1 - c)^t = 1 - tc + \frac{t(t-1)(1-s)^{t-2}}{2} c^2 \quad (102)$$

for some $s \in (0, c)$. Consider the function

$$g(t, \lambda, y) := 1 + \lambda - \lambda \cdot (1 + y)^t - (1 - \lambda y)^t. \quad (103)$$

Equation (102) yields that

$$\begin{aligned} g(t, \lambda_i, y_i) &= 1 + \lambda_i - \lambda_i \cdot (1 + y_i)^t - \left(1 - t\lambda_i y_i + \frac{t(t-1)(1-s_i)^{t-2}}{2} \lambda_i^2 y_i^2 \right) \\ &= \lambda_i \left(1 - (1 + y_i)^t + ty - \frac{t(t-1)(1-s_i)^{t-2}}{2} \lambda_i y_i^2 \right) \end{aligned} \quad (104)$$

for every index i and some $s_i \in (0, \lambda_i y_i)$. We conclude that

$$\begin{aligned} \lim_{i \rightarrow \infty} f_{w_i}(\lambda_i, y_i) &= \lim_{i \rightarrow \infty} \frac{g(1 + w_i, \lambda_i, y_i)}{g(2 + w_i, \lambda_i, y_i)} \\ &= \lim_{i \rightarrow \infty} \frac{1 - (1 + y_i)^{1+w_i} + (1 + w_i)y_i - \frac{(1+w_i)w_i(1-s_i)^{w_i-1}}{2} \lambda_i y_i^2}{1 - (1 + y_i)^{2+w_i} + (2 + w_i)y_i - \frac{(2+w_i)(1+w_i)(1-s_i)^{w_i}}{2} \lambda_i y_i^2} \\ &= \lim_{i \rightarrow \infty} \frac{\frac{1}{(1+y_i)^{2+w_i}} - \frac{(1+y_i)^{1+w_i}}{(1+y_i)^{2+w_i}} + \frac{(1+w_i)y_i}{(1+y_i)^{2+w_i}} - \frac{(1+w_i)w_i(1-s_i)^{w_i-1} \lambda_i y_i^2}{2(1+y_i)^{2+w_i}}}{\frac{1}{(1+y_i)^{2+w_i}} - 1 + \frac{(2+w_i)y_i}{(1+y_i)^{2+w_i}} - \frac{(2+w_i)(1+w_i)(1-s_i)^{w_i} \lambda_i y_i^2}{2(1+y_i)^{2+w_i}}} \\ &= 0. \end{aligned}$$

$\lambda^* = \infty$: Note that the assumption $\lambda^* = \infty$ yields that $y^* = 0$. For $c \in [0, 1)$, the Taylor expansion with Lagrange remainder over the interval $[0, c]$ yields that

$$(1 - c)^t = 1 - tc + \frac{t(t-1)}{2} c^2 - \frac{t(t-1)(t-2)(1-s)^{t-3}}{6} c^3, \quad (105)$$

for some $s \in (0, c)$, and

$$(1 + c)^t = 1 + tc + \frac{t(t-1)}{2} c^2 + \frac{t(t-1)(t-2)(1+s')^{t-3}}{6} c^3, \quad (106)$$

for some $s' \in (0, c)$.

Applying Equations (105) and (106) for the function g of Equation (103) yields that

$$\begin{aligned}
& g(t, \lambda_i, y_i) \tag{107} \\
&= \tilde{g}(t, \lambda_i, y_i, s_i, s'_i) \\
&:= 1 + \lambda_i - \lambda_i \left(1 + ty + \frac{t(t-1)}{2} y_i^2 + \frac{t(t-1)(t-2)(1+s'_i)^{t-3}}{6} y_i^3 \right) \\
&\quad - \left(1 - t\lambda_i y_i + \frac{t(t-1)}{2} \lambda_i^2 y_i^2 + \frac{t(t-1)(t-2)(1-s_i)^{t-3}}{6} \lambda_i^3 y_i^3 \right) \\
&= -\frac{\lambda_i^2 y_i^2}{6} \left(\frac{3t(t-1)}{\lambda_i} + \frac{t(t-1)(t-2)(1+s'_i)^{t-3} y_i}{\lambda_i} + 3t(t-1) + t(t-1)(t-2)(1-s_i)^{t-3} \lambda_i y_i \right)
\end{aligned}$$

for large enough index i and some $s_i \in (0, \lambda_i y_i)$ and $s'_i \in (0, y_i)$. We conclude that

$$\begin{aligned}
& \lim_{i \rightarrow \infty} f_{w_i}(\lambda_i, y_i) \\
&= \lim_{i \rightarrow \infty} \frac{g(1+w_i, \lambda_i, y_i)}{g(2+w_i, \lambda_i, y_i)} \\
&= \lim_{i \rightarrow \infty} \frac{\tilde{g}(1+w_i, \lambda_i, y_i, s_i, s'_i)}{\tilde{g}(2+w_i, \lambda_i, y_i, s_i, s'_i)} \\
&= \lim_{i \rightarrow \infty} \frac{\frac{3(1+w_i)w_i}{\lambda_i} + \frac{(1+w_i)w_i(w_i-1)(1+s'_i)^{w_i-1} y_i}{\lambda_i} + 3(1+w_i)w_i + (1+w_i)w_i(w_i-1)(1-s_i)^{w_i-2} \lambda_i y_i}{\frac{3(2+w_i)(1+w_i)}{\lambda_i} + \frac{(2+w_i)(1+w_i)w_i(1+s')^{w_i-1} y_i}{\lambda_i} + 3(2+w_i)(1+w_i) + (2+w_i)(1+w_i)w_i(1-s)^{w_i-1} \lambda_i y_i} \\
&= \frac{0}{6} = 0,
\end{aligned}$$

where the next-to-last equality holds since $\lambda_i y_i \leq 1$ for every i , and hence the last term of the numerator and denominator goes to 0 when $i \rightarrow \infty$.

$\lambda^*, y^* > 0$: It holds that

$$\begin{aligned}
\lim_{i \rightarrow \infty} f_{w_i}(\lambda_i, y_i) &= \lim_{i \rightarrow \infty} \frac{1 + \lambda_i - \lambda_i \cdot (1 + y_i)^{1+w_i} - (1 - \lambda_i y_i)^{1+w_i}}{1 + \lambda_i - \lambda_i \cdot (1 + y_i)^{2+w_i} - (1 - \lambda_i y_i)^{2+w_i}} \\
&= \frac{1 + \lambda^* - \lambda^*(1 + y^*) - (1 - \lambda^* y^*)}{1 + \lambda^* - \lambda^*(1 + y^*)^2 - (1 - \lambda^* y^*)^2} \\
&= 0.
\end{aligned}$$

$\lambda^* = 0$ and $y^* > 0$: Equations (102) and (104) yield that

$$\begin{aligned}
\lim_{i \rightarrow \infty} f_{w_i}(\lambda_i, y_i) &= \lim_{i \rightarrow \infty} \frac{1 - (1 + y_i)^{1+w_i} + (1 + w_i)y_i - \frac{(1+w_i)w_i(1-s_i)^{w_i-1}}{2} \lambda_i y_i^2}{1 - (1 + y_i)^{2+w_i} + (2 + w_i)y_i - \frac{(2+w_i)(1+w_i)(1-s_i)^{w_i}}{2} \lambda_i y_i^2} \\
&= \frac{1 - (1 + y^*) + y^*}{1 - (1 + y^*)^2 + 2y^*} \\
&= 0.
\end{aligned}$$

$y^* = 0$: Rearranging Equation (107) yields that the following holds for large enough index i :

$$\begin{aligned}
 g(t, \lambda_i, y_i) & & (108) \\
 &= \tilde{g}(t, \lambda_i, y_i, s_i, s'_i) \\
 &= -\frac{\lambda_i y_i^2}{6} (3t(t-1) + t(t-1)(t-2)(1+s'_i)^{t-3} y_i + 3t(t-1)\lambda_i + t(t-1)(t-2)(1-s_i)^{t-3} \lambda_i^2 y_i)
 \end{aligned}$$

for some $s_i \in (0, \lambda_i y_i)$ and $s_i \in (0, y_i)$. Given, this formulation it is easy to see that

$$\begin{aligned}
 \lim_{i \rightarrow \infty} f_{w_i}(\lambda_i, y_i) &= \lim_{i \rightarrow \infty} \frac{\tilde{g}(1+w_i, \lambda_i, y_i, s_i, s'_i)}{\tilde{g}(2+w_i, \lambda_i, y_i, s_i, s'_i)} \\
 &= \frac{0}{6+6\lambda^*} \\
 &= 0.
 \end{aligned}$$

The above holds since every term in the numerator goes to 0 and the term $3(2+w_i)(1+w_i)$ in the denominator goes to 6.

This concludes the case analysis, and thus the proof of the claim. □