The Price of Agreement in Secure Multiparty Computation

Ran Cohen* Pouyan Forghani† Juan Garay‡ Vassilis Zikas§

November 3, 2026

Abstract

Secure multi-party computation (MPC) cannot guarantee output delivery, in general, in the dishonest-majority setting. Security of MPC is therefore relaxed to security with abort, where the adversary may prevent honest parties from obtaining their output. The standard notion of *unanimous abort* (UA-MPC) ensures that either all honest parties obtain their output or all of them abort. Goldwasser and Lindell [JCRYPT '05] introduced the relaxed notion of *selective abort*, allowing some honest parties to abort and others to obtain their output. While two-round SA-MPC can be achieved over point-to-point channels, two-round UA-MPC necessitates the second round to be over broadcast. The price of agreement in dishonest-majority MPC is thus tied to the round complexity of broadcast with unanimous abort (UA-broadcast).

The round complexity of broadcast has been extensively studied for more than four decades. Garay, Katz, Koo, and Ostrovsky [FOCS '07] (GKKO) proved that the expected round complexity of broadcast protocols tolerating t corruptions is at least $2n/(n-t)-1$. Their argument, however, relies on all honest parties obtaining an honest sender's input, and does not extend to UA-broadcast. Recently, Cohen *et al.* [TCC '26] showed that UA-MPC is possible in an expected constant number of rounds when a constant fraction of the parties is honest, and asked whether this extends to $n - o(n)$ corruptions.

In this work we answer this question in the negative: Agreement alone already incurs the GKKO barrier. Concretely, we prove that for $t \leq n-2$, UA-broadcast—and hence UA-MPC—requires at least $t/(2(n-t))$ rounds in expectation, matching the asymptotic bound for broadcast without abort. As a corollary, expected-constant-round UA-MPC is possible if and only if a constant fraction of the parties is honest.

Next, we consider UA-broadcast protocols that terminate in an *a-priori* known round, and show that the bound that the argument of GKKO gives for broadcast holds for them exactly. Combined with the broadcast protocol of Dolev and Strong [SICOMP '83], our bound is tight for arbitrarily many corruptions.

*Efi Arazi School of Computer Science, Reichman University. E-mail: cohenran@runi.ac.il.

†School of Cybersecurity and Privacy, Georgia Institute of Technology. E-mail: pforghani3@gatech.edu.

‡Department of Computer Science and Engineering, Texas A&M University. E-mail: garay@tamu.edu.

§School of Cybersecurity and Privacy, Georgia Institute of Technology. E-mail: vzikas@gatech.edu.

Contents

1 Introduction **1**
1.1 Our Contributions 2
1.2 Overview of Our Techniques 4

2 Preliminaries **8**
2.1 Model 9
2.2 Broadcast with Unanimous Abort 9

3 Lower Bound for Broadcast with Unanimous Abort **10**
3.1 Executions on a Line 10
3.2 The Lower Bound 15

4 Lower Bound for Deterministic Broadcast with Unanimous Abort **20**

AI Methodology **25**

Bibliography **25**

A The Lower Bound against Oblivious Adversaries **28**

1 Introduction

Secure multi-party computation (MPC) allows n parties to jointly compute a function over their private inputs, even in the presence of t corrupted parties. Ideally, MPC protocols guarantee output delivery (GD), meaning that every honest party receives its output regardless of the adversary's strategy [BGW88, CCD88, RB89]. In particular, guaranteed output delivery implies a form of *agreement*, as all honest parties obtain correlated outputs consistent with the same inputs.

Without an honest majority, however, guaranteed output delivery is impossible in general [Cle86]. Consequently, Goldreich, Micali, and Wigderson [GMW87] relaxed the security of MPC to *security with abort*, by allowing the adversary to first learn the output values of all corrupted parties and then decide whether to allow honest parties to obtain their output or not. Although security with abort gives up on guaranteed output delivery, standard definitions maintain agreement by ensuring *unanimous abort* (UA): either all honest parties receive their output or all abort. Every function can be computed with UA against any number of corruptions in two rounds, provided the first round uses point-to-point channels and the second uses a broadcast (BC) channel [CGZ20].

Goldwasser and Lindell [GL05] introduced a further relaxation called *security with selective abort* (SA), which dispenses with agreement entirely by allowing the adversary to selectively choose which honest parties obtain output and which abort. Security with selective abort removes the need for broadcast altogether, and two rounds over point-to-point channels are sufficient to compute any function [CGZ20].

The focus of this work is the *price of agreement* in MPC: the number of point-to-point rounds needed to achieve MPC with unanimous abort (UA-MPC) beyond the two rounds needed to achieve MPC with selective abort (SA-MPC). This is tied to the round complexity of broadcast with unanimous abort (UA-BC), in which honest parties either agree on the message sent by an honest sender or all identify an attack and abort.

The round complexity of broadcast. Broadcast protocols enable a sender to distribute its message to a set of parties such that all honest parties reach agreement on a value which equals the sender's message as long as the sender is honest [PSL80, LSP82]. Dolev and Strong [DS83] constructed a broadcast protocol secure against any number of corruptions that terminates in $t+1$ rounds. This protocol is deterministic and perfectly secure, given black-box access to ideal signatures [CDL+25], and ensures that all honest parties terminate together in an *a-priori* known round (round $t+1$). Matching lower bounds show that $t+1$ rounds are necessary for such deterministic protocols [FL82, DLM82, DS83].

Ben-Or [Ben83] and Rabin [Rab83] showed that randomized broadcast protocols can circumvent this bound. Following a fruitful line of work, broadcast protocols with an expected constant number of rounds were constructed by Feldman and Micali for $t < n/3$ [FM88], and by Katz and Koo for $t < n/2$ in the PKI model [KK06]. Expected-constant-round protocols can no longer ensure simultaneous termination, and present challenges when composing such protocols in a larger MPC protocol [BE03, LLR02, CCGZ19, CCGZ21, CFG+23]. Randomized broadcast protocols with simultaneous termination that terminate in R rounds err with probability at least $(c \cdot R)^{-R}$ for some constant c [KY84, CMS89]; a protocol matching this bound was introduced by Ghinea, Goyal, and Liu-Zhang [GGLZ22].

In the dishonest-majority setting, broadcast with expected constant rounds for $t = n/2 + O(1)$ was constructed by Garay, Katz, Koo, and Ostrovsky [GKKO07] and by Fitzi and Nielsen [FN09]. Garay *et al.* proved a lower bound showing that every randomized broadcast protocol requires $2n/(n-t) - 1$ rounds in expectation. Together with [DS83], this determines the round complexity of broadcast up to a constant factor for $t = n - O(1)$. When considering a constant fraction of

honest parties, Chan, Pass, and Shi constructed a protocol with $\mathsf{polylog}(n)$ rounds and simultaneous termination [CPS20], and Wan, Xiao, Shi, and Devadas a protocol with expected constant rounds in a PKI for verifiable random functions [WXSD25].

Broadcast and MPC with abort. Relaxing output delivery changes the picture. With unanimous abort the honest parties agree on whether to abort, and with identifiable abort (IA) [IOZ14] they additionally agree on a corrupted party to blame. UA-broadcast was studied more than two decades ago as detectable broadcast [FGM01, FGMvR02], and Fitzi *et al.* [FGH+02] achieved it for every $t < n$ in $t + 3$ rounds without any setup—which broadcast cannot do once $t \geq n/3$. Giving up agreement as well removes almost all rounds: a simple two-round echo replaces every broadcast round [GL05], and two rounds over point-to-point channels suffice for MPC with selective abort, assuming two-round oblivious transfer [CGZ20]. Recently, Cohen *et al.* [CCF+26] showed that UA-MPC runs in an expected constant number of rounds whenever a constant fraction of the parties is honest, by letting one party broadcast the signed messages of all parties within a single instance of the protocol of [WXSD25]. They also showed that identifiable abort requires $\Omega(n/(n-t))$ rounds, by extending the argument of [GKKO07], and that UA-broadcast with a fixed round bound requires a super-constant number of rounds against $t = \omega(1)$ corruptions, by extending [CMS89, KY84]. This leaves the case of few honest parties open, and they asked the following question:

Is expected-constant-round UA-MPC possible in the presence of $n - o(n)$ corruptions?

Answering this question requires a lower bound on the expected round complexity of UA-broadcast that is super-constant when $n - t = o(n)$, and none of the known techniques provides one. The chain argument of Dolev and Strong [DS83, Du89] applies to UA-broadcast—the ends of its chain are all-honest executions—but only to protocols with perfect guarantees, for which it yields $\min\{t + 1, n - 1\}$ rounds. Perfect guarantees exclude randomized protocols and every protocol that relies on computational assumptions, including that of Dolev and Strong itself once its signatures are instantiated. Its randomized versions [CMS89, KY84], which [CCF+26] extended to UA-broadcast, allow negligible error but apply only to protocols with a fixed round bound, and say nothing about the expected number of rounds. The argument of [GKKO07] does bound the expected number of rounds, but it relies on a single execution in which an honest sender at one end of a line of parties forces the output; under unanimous abort, this execution may simply end with every party aborting. Its extension to identifiable abort [CCF+26] rules out such aborts, since all honest parties would have to blame the same party along the whole line, including where that party is honest. Under unanimous abort, however, the honest parties abort without blaming anyone.

In Table 1 and in our theorems, we combine the shorthands GD, IA, UA and SA with BC and MPC in the same way, so that, for instance, IA-BC is broadcast with identifiable abort and SA-MPC is MPC with selective abort.

1.1 Our Contributions

In this work we answer the above question in the negative: agreement alone already incurs the barrier of Garay *et al.* [GKKO07]. Stated differently, with a constant number of honest parties neither randomization nor abort saves more than a constant factor over the $t + 1$ rounds of Dolev and Strong [DS83]. Table 1 places our results among the known bounds.

Theorem 1.1. *Let $t \leq n - 2$ and let π be any t-secure UA-BC protocol against rushing omission adversaries, then its expected round complexity is at least $\frac{t}{2(n-t)}$.*

		$t \le n-2$	$t \le (1-\varepsilon)n$	$t = n-c$		$t \ge n-2$	
		LB (Exp.)	UB (Exp.)	LB (Exp.)	UB	LB	UB
BC	GD	$\frac{2n}{n-t}-1$ [GKKO07]	$O(1)$ [WXSD25]	$\frac{2t}{c}+1$ [GKKO07]	$t+1$ [DS83]	$n-1$ [GKKO07]	$n-1$ [DS83]
	IA	$\frac{2n}{n-t}-1$ [CCF+26]	↓	$\frac{2t}{c}+1$ [CCF+26]	↓	$n-1$ [CCF+26]	↓
	UA	$\boldsymbol{\frac{t}{2(n-t)}}$ **Thm. 1.1**	↓	$\boldsymbol{\frac{t}{2c}}$ **Thm. 1.1**	↓	$\boldsymbol{n-1}$ **Thm. 1.2**	↓
	SA	← 2 [GL05] →					
MPC	IA	$\Omega\left(\frac{n}{n-t}\right)$ [CCF+26]	$O(1)$ [CEG26, CGZ20]	$\Omega(t)$ [CCF+26]	$O(t)$ [CGZ20, DS83]	$\Omega(n)$ [CCF+26]	$O(n)$ [CGZ20, DS83]
	UA	$\boldsymbol{\Omega\left(\frac{t}{n-t}\right)}$ **Thm. 1.1**	$O(1)$ [CCF+26]	$\boldsymbol{\Omega(t)}$ **Thm. 1.1**	↓	$\boldsymbol{\Omega(n)}$ **Thm. 1.2**	↓
	SA	← 2 [CGZ20] →					

Table 1: Round complexity over point-to-point channels against t static corruptions, for a constant $\varepsilon > 0$ for broadcast and MPC. GD, IA, UA and SA stand for guaranteed output delivery and for identifiable, unanimous and selective abort, and LB and UB for lower and upper bounds. Columns marked Exp. bound the expected round complexity, and the others the round complexity of protocols with a fixed round bound. In the column $t = n - c$, $c \ge 2$ is a constant, and an arrow repeats an entry in the cells it points to. A protocol secure against t corruptions is also secure against fewer, so every lower bound for t also holds for larger t, and the expected bounds of the column $t = n - c$ for $c = 2$ hold for $t = n - 1$ as well. In each part of the table, the rows relax output delivery further, and only selective abort gives up agreement. The upper bounds hold under the setup and assumptions of the cited works, and our lower bounds hold for every setup that is independent of the inputs.

The bound holds against a static rushing *omission* adversary, whose corrupted parties follow the protocol except that they withhold or ignore messages [PT86, Had84], as in [GKKO07]. It also holds for every setup that is sampled independently of the inputs, e.g., a PKI, a common reference string, or the honestly generated keys for verifiable random functions used by [WXSD25, CCF+26].[1] This is asymptotically the same as the bound $2n/(n-t)-1$ of [GKKO07] for broadcast,[2] and it is asymptotically tight when a constant number of parties is honest, as [DS83] then runs in $t+1 = \Theta(t/(n-t))$ rounds (Section 3.2).

A UA-MPC protocol for the function that hands the sender's input to every party is a UA-broadcast protocol, so all our bounds apply to UA-MPC as well. As a corollary, together with the protocol of [CCF+26]—which assumes oblivious transfer, verifiable random functions, and a PKI—Theorem 1.1 shows that expected-constant-round UA-MPC over point-to-point channels is possible exactly when a constant fraction of the parties is honest, and that the protocol of [CCF+26]

[1] A rushing adversary sees the messages that the honest parties send to corrupted parties in a round before deciding what the corrupted parties do in that round. This is the standard adversary of the synchronous model, and the protocol of [CCF+26] is secure against it. The expected round complexity is the supremum, over all adversaries and inputs, of the expected round in which the last honest party terminates (cf. Section 2.2).

[2] As the proof in Section 3.2 notes, every UA-BC protocol takes at least one round in expectation, since a party other than the sender learns nothing about the sender's input before the first round. So the bound is in fact $\max\{1, t/(2(n-t))\}$, which matches $\Omega(n/(n-t))$ also for $t = o(n)$.

has optimal resilience. Conversely, UA-MPC can be built from a constant number of rounds and a single instance of UA-broadcast [CGZ20, CCF+26], so under the assumptions of these works the round complexities of the two differ by at most an additive constant.

Our second result concerns protocols with a fixed round bound (aka deterministic termination), and holds already against the weakest omission adversary, namely an *oblivious* one, which fixes the corrupted parties and the messages they withhold or ignore before the setup and the parties' randomness are sampled.

Theorem 1.2. *Let $t < n$. For any R-round t-secure UA-BC protocol, R is at least* $\min\{n-1, \lceil \frac{n}{n-t} \rceil + \lfloor \frac{n}{n-t} \rfloor - 1\}$*, which is at least* $\frac{2n}{n-t} - 2$ *if* $t \leq n-2$.

The bound of Theorem 1.2 is the distance between the ends of the longest line on which the argument of [GKKO07] runs. It is thus exactly what their argument gives for broadcast, and equals their bound $2n/(n-t)-1$ when $(n-t)/2$ divides n. For $t \geq n-2$ it equals $n-1$, which is exactly the round complexity of [DS83] for $n-2$ corruptions; the same protocol also tolerates $n-1$ corruptions if the sender outputs its input at the start (Section 4). Hence Theorem 1.2 is exactly tight when at most two parties are honest, and asymptotically tight for any constant number of honest parties, as the bound and $t+1$ are then both $\Theta(n)$. We stress that it holds for all protocols with a fixed round bound—even randomized ones with negligible error that may abort and whose parties may terminate in different rounds—whereas it is matched by a protocol that is deterministic, never aborts, and has simultaneous termination. Since [GKKO07] already covers protocols that never abort and [DS83] protocols with perfect guarantees, the new case is that of protocols that may abort and have negligible error. In particular, when at most two parties are honest, Theorem 1.2 shows that the protocol of Dolev and Strong, which has a negligible probability of error once its signatures are instantiated, is round-optimal even among protocols that may abort, whereas their own lower bound, proved for perfect guarantees, does not apply to it. Together with [CMS89, KY84], which also apply to UA-broadcast [CCF+26], each of the classical lower bounds for broadcast with a fixed round bound thus carries over to UA-broadcast.

In a nutshell, wherever the round complexity of broadcast is known—with a constant fraction or a constant number of honest parties—the price of agreement equals it up to a constant factor. In between, Theorem 1.1 matches the lower bound of [GKKO07] for broadcast up to a constant factor, so giving up output delivery while keeping agreement can save more than a constant factor only if broadcast itself requires asymptotically more than $n/(n-t)$ rounds.

1.2 Overview of Our Techniques

Both our results start from the line argument of [GKKO07]. We first recall it and explain why it breaks under unanimous abort, then describe the chains of executions on which all our proofs rest, and finally explain how each result instantiates them, in the order in which the techniques build on each other.

The line argument. At a high level, [GKKO07] place the parties in sets along a line, with the sender's set at one end, such that any two adjacent sets jointly contain at least $n-t$ parties. We use the longest such line: it starts with the sender alone, then alternates sets of $n-t-1$ parties and single parties, and ends with the remaining parties, if any (Figure 3 in Section 3.1). For $t \leq n-2$ we write d for the distance between its ends, which is $\lceil n/(n-t) \rceil + \lfloor n/(n-t) \rfloor - 1$ and thus the bound of Theorem 1.2; e.g., for $t = n-2$ every set is a single party and $d = n-1$. We call a message between two non-adjacent sets a *cross message*. In the *line execution*, all parties follow

the protocol but all cross messages are dropped, so every set hears only from itself and its two neighbors.

The key observation is that the line execution is the execution under many attacks at once. Indeed, for every pair of adjacent sets, it is exactly what happens when an adversary corrupts everyone outside the pair and has them drop the cross messages—a legal attack, as the pair contains at least $n - t$ parties. Hence every guarantee that the protocol gives against any of these attacks holds in the line execution simultaneously; in particular, agreement within each pair, chained along the overlapping pairs, gives all parties a common output. The same holds for every execution that drops only cross messages. Whereas [GKKO07] use this for the line execution alone, we use it for every execution we consider, which is what lets us compare executions with the same inputs and randomness.

For broadcast, the two ends of the line execution give the bound. At the sender's end, the sender is honest in the attack that keeps its own set and the next one honest, so output delivery makes the common output the sender's input. At the other end, information moves only from a set to its neighbors, one set per round (Lemma 3.1), so the far party learns nothing about the sender's input before round d and can output it earlier only by guessing, which fails with probability $1/2$.

Starting from the all-honest execution. Unfortunately, unanimous abort breaks this argument at the sender's end. An honest sender now only guarantees that the honest parties output its input *or abort*, and since in the line execution every pair of adjacent sets is under attack, all parties may abort. Indeed, a protocol in which a party aborts as soon as a message goes missing does exactly this, and its far party never needs to learn the input. The only execution in which unanimous abort forces the output to be the sender's input is the all-honest one, and all our proofs start there.

To get from the all-honest execution to one in which the far party is blind, we use a chain of executions, each dropping a few more cross messages than the previous one. As only cross messages are dropped, the above observation gives every execution of the chain a common output. We call two consecutive executions of the chain a *link*. They differ only in messages of a single round into some of the sets, and both drop all cross messages of the later rounds, so their difference spreads by one set per round (Figure 1). A party far from the changed sets—the *witness* of the link—sees the same thing in both executions until the difference arrives. If it outputs before then, it outputs the same value in both, and by agreement the two executions have the same common output. Hence, if every witness outputs in time, the common output stays the sender's input along the whole chain, up to its end, where the far party cannot know the input in time. In other words, some witness, or the far party at the end, must be slow. Each of our results builds its chain such that every witness is far from the change it watches, and turns these slow parties into a bound in its own way. We stress that the line, the attack on each adjacent pair, the common output of an execution, and the blindness of the far party are from [GKKO07]. What is new is a chain from the all-honest execution to one in which the far party is blind, whose witnesses notice each change only after it has spread along the line to them, as well as the adversaries of Theorem 1.1 and the rotating cycle of Theorem 1.2.

The bound with probabilistic termination. Theorem 1.1 uses the simplest chain, on the line itself. The crux is that an output in the all-honest execution is a commitment: once a party has output the sender's input, an attack that starts later and leaves this party honest can no longer make the honest parties abort, so from then on the protocol must deliver the input to every party, as a broadcast protocol would. Our adversary exploits this commitment.

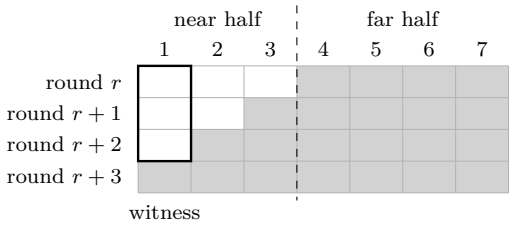


Figure 1: One link on a line of seven sets, with the sender alone in set 1 and the near half consisting of sets 1–3. The two executions of the link differ only in the messages of round r into the far half, and both drop all cross messages of later rounds. A shaded cell means that a party in that set may see something different at the end of that round in the two executions. The difference spreads by one set per round, so the sender—the witness of this link—sees the same thing in both executions through round $r+2$, and if it outputs by then, it outputs the same value in both.

We say that an execution *cuts* the line at a round if it drops all cross messages from that round on; the line execution cuts at round 1 and the all-honest execution does not cut at all. The chain starts from the execution that cuts the line at round $L = \lfloor n/(n-t) \rfloor$, which agrees with the all-honest execution through round $L-1$, and moves the cut back one round at a time, to rounds $L-1, \ldots, 1$. Each move takes two links: the first drops the cross messages of that round into the far half of the line, with the sender as the witness, and the second drops those into the near half—the one containing the sender—with the far party as the witness (Figure 1; Figure 6 in Section 3.1 shows the whole chain for a line of seven sets). The cut moves backward because a difference spreads one set per round only once the later rounds are cut, and each round is split into halves because dropping a whole round at once would change what every set sees. Every witness is at distance at least L from the half it watches, so it sees the same thing in both executions of its link through round L.

A natural idea would be to let an oblivious adversary realize an execution of the chain with a slow witness. This, however, fails, as a protocol may spread its slowness evenly, with a different witness slow on different coins. An oblivious adversary fixes its omissions before the execution starts, so it realizes a single execution of the chain, which may rarely have a slow witness; against such adversaries our argument yields only $\Omega(\sqrt{n/(n-t)})$ rounds (Appendix A). Instead, we use a rushing omission adversary that decides *during* the execution, from what its parties have seen so far, whether to cut the line at the current round. Its corrupted parties follow the protocol and never send a wrong message, and it thus realizes the execution that cuts at a round of its choosing, or the all-honest execution if it never cuts.

The challenge is to recognize the right round in time. To this end, we watch two parties: the sender's neighbor in the second set, and the far party. (The adversary has to corrupt the sender to know its input, which is why we watch the neighbor rather than the sender itself.) In an execution, the *status* of a party is yes if it outputs the sender's input by round $L-1$, no if it outputs something else by then, and slow otherwise.[3] Consider the list of statuses of the far party when the line is cut at round 1, of the neighbor when it is cut at round 2, of the far party when it is cut at round 2, of the neighbor when it is cut at round 3, and so on, up to the far party when it is cut at round L. Two consecutive statuses in this list are those of the two parties in a single execution. This is clear for the neighbor and the far party when the line is cut at the same round. The far party when it is cut at round ρ and the neighbor when it is cut at round $\rho+1$ have the same statuses in the execution between the two links that move the cut from $\rho+1$ to ρ, since each of them is at distance

[3]The deadline is one round earlier than L because the neighbor is one set closer to the far half than the sender.

at least $L-1$ from the half in which its execution differs from that one. So agreement never puts a yes next to a no. The list starts with the far party in the line execution, which is almost never yes: the far party is blind there, so it outputs the input in time exactly as often as the other bit, and validity for the honest sender of one of the attacks on the line execution rules out the other bit. The list ends with the far party in an execution that agrees with the all-honest one through round $L-1$, which is yes or slow by validity. Hence, either the list contains no yes and ends with slow, or its first yes comes right after a slow.

Two adversaries find the slow party. The first keeps the last two sets honest—so the far party is honest—and corrupts everyone else, including the sender. In every round, it computes the status that the neighbor would have if the line were cut at the next round, and cuts at the current round as soon as this status is yes, i.e., just before the neighbor would commit to the sender's input. If the first yes in the list is the neighbor's, the slow right before it is the status of the honest far party in exactly the execution that this adversary realizes; if the list has no yes, the adversary never cuts and the far party is slow in the all-honest execution. The adversary can compute the neighbor's status because the neighbor and everyone near it are corrupted, and nothing that the honest parties at the far end send after the cut reaches the neighbor by round $L-1$ (Claim 3.3). It does, however, need the honest messages of the current round before deciding whether to cut at that round, which is the one place where rushing is used. The second adversary, the mirror image of the first, keeps the second and third sets honest—so the neighbor is honest—and cuts at the first round at which the far party would have status yes if the line were cut there. It covers the case in which the first yes is the far party's, and needs no rushing.

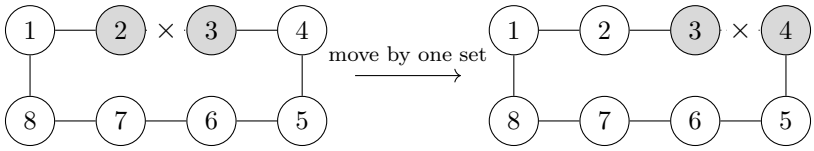
Thus, on almost every run, one of the two adversaries realizes an execution in which an honest party has not output by round $L-1$. The probabilities of this event under the two adversaries add up to almost 1, so one of them is at least about $1/2$, and the expected output round of that party is at least about $L/2 \geq (n/(n-t)-1)/2 = t/(2(n-t))$. As the all-honest execution is just the first adversary's execution when it never cuts, no assumption on the all-honest execution is needed. Our attack is in the spirit of Cleve's attack on coin flipping [Cle86], which aborts as soon as the backup value of the honest parties is the one the adversary wants; here, the backup value is the output that the sender's neighbor would have after a cut one round later.

The bound with deterministic termination. With a fixed round bound R, every party outputs by round R in every execution of a chain, except with negligible probability. A link then keeps the common output whenever its witness sees the same thing in both executions through round R, and loses only the negligible probability with which the guarantees fail, so the chain may have polynomially many links. If a link changes messages of round r and both its executions drop all cross messages of later rounds, the witness sees the same thing through round R when the changed sets are at distance at least $R-r+1$ from it. The chain of Theorem 1.1, however, falls short: in its link that cuts round 1 into the far half, the witness is the sender, at distance L from the change, so it sees the same thing only through round L, and the chain yields only $R \geq L+1$, about half of d.

Suppose that $R \leq d-1$. We let the chain end in an execution in which the far party knows nothing about the sender's input through round R. As only the sender's messages depend on its input, it suffices to drop all cross messages from round 2 on, together with the sender's round-1 messages into every set within distance $R-1$ of the far end, since from any other set the input reaches the far end only after round R (Figure 8(b) in Section 4). The difficulty is to reach this execution by links. A link that drops messages of round r into a set needs a witness at distance at least $R-r+1$ from that set, which for the early rounds is almost the length d of the line. The set must then sit at one end of the line and the witness at the other, whereas the cross messages of

each round have to be dropped into every set of the line, and the sender's round-1 messages into up to R sets—and a line has only two ends.

We therefore place the sets on a cycle, on which every set can be an end. The sets are those of the line, except that if $n - t$ does not divide n its last two sets are merged, so that the last set and the sender's set also jointly contain at least $n - t$ parties. *Opening* the cycle between two adjacent sets yields a line with these two sets as its ends (Figure 2), and moving the opening by one set takes only links whose changed set and witness are the two ends of the current line (Claim 4.1). The move restores the messages across the old opening round by round from early to late, and then drops those across the new opening from late to early, so that every link finds the cross messages of all later rounds dropped, as it requires. The chain cuts the rounds from R down to 2, as in Theorem 1.1, but set by set, so for each round it moves the opening once around the cycle and drops the cross messages of that round into each set while that set is an end. For round 1 it moves the opening to each set near the far end in turn and drops the sender's messages into it, and if a set was merged, a final phase splits it again. As the chain has polynomially many links, all guarantees hold along it except with negligible probability, and the common output stays the sender's input. Hence the far party outputs the sender's input by round R except with negligible probability, although it knows nothing about this input through round R and can thus output it in time with probability at most $1/2$.



(a) the line $3, 4, \ldots, 8, 1, 2$ (b) the line $4, \ldots, 8, 1, 2, 3$

Figure 2: Opening a cycle of eight sets, with the sender alone in set 1. Dropping the messages across the opening, marked ×, and all messages between sets that are not adjacent on the cycle leaves a line whose ends (shaded) are the two sets next to the opening. Moving the opening by one set, from (a) to (b), changes one end of the line, so moving it around the cycle makes every set an end in turn.

Organization of the paper. In Section 2 we define the model and UA-broadcast. In Section 3 we prove the bound for probabilistic termination, and in Section 4 the bound for deterministic termination, i.e., with a fixed round bound. In Appendix A we prove the bound against oblivious omission adversaries.

2 Preliminaries

Conventions. We consider n parties, and we identify each party P_i with $i \in [n]$. We write κ for the security parameter, and we assume that $2 \leq n = \mathsf{poly}(\kappa)$ and that honest parties run in polynomial time. Our proofs lose only negligible terms, which the bounds absorb for all sufficiently large κ, and for readability their statements leave this condition on κ implicit. At most $t < n$ parties are corrupted, and we write $h \triangleq n - t$ for the minimum number of honest parties. Throughout, we let

$$d \triangleq \min\{n - 1, \lceil n/h \rceil + \lfloor n/h \rfloor - 1\}.$$

This is the length of the longest line of sets on which our lower bounds run (Section 3.1). If $h \geq 2$ and $h/2$ divides n, it equals $2n/h - 1$, the bound that Garay *et al.* [GKKO07] prove in this case.

The term $n-1$ matters only for $h=1$, since for $h \geq 2$ the second term is at most the first.

2.1 Model

Network. The parties communicate over authenticated point-to-point channels in the synchronous model. That is, the execution proceeds in rounds, and a message that a party sends in a round is delivered to its recipient at the end of that round. So the messages that an honest party sends in a round depend only on its input, its setup, its random tape and the messages it received in earlier rounds. Before the execution, the parties may receive a setup that is sampled independently of their inputs, such as a PKI or a common reference string.

Adversaries. The adversary is probabilistic polynomial-time (PPT) and static, that is, it corrupts a set of at most t parties before the execution starts. We call such an adversary a *t-adversary*. Our definitions require security against every t-adversary, whose corrupted parties may deviate arbitrarily from the protocol. Our lower bounds, however, use only *omission adversaries*, whose corrupted parties follow the protocol, except that they may withhold messages they should send and ignore messages they receive [PT86, Had84, ZHM09]. At the end of the execution, such an adversary outputs the outputs of the corrupted parties. We consider two kinds of omission adversaries.

- A *rushing omission adversary* chooses its omissions during the execution, as in the standard omission model. It is *rushing*, that is, in every round it learns the messages that the honest parties send to corrupted parties in that round before it decides which messages the corrupted parties withhold or ignore in that round. Theorem 1.1 is stated against such adversaries.
- An *oblivious omission adversary* is weaker. It is not rushing, and it fixes the corrupted parties and the messages they withhold or ignore, each given by its sender, receiver and round, before the setup, the inputs and the random tapes of the parties are sampled, like the oblivious adversaries of randomized consensus [Asp12]. A rushing omission adversary can act as one by ignoring what it learns, so security against rushing omission adversaries implies security against oblivious ones. The proofs of Theorems 1.2 and A.1 use only such adversaries.

2.2 Broadcast with Unanimous Abort

Definition 2.1 (UA-BC)**.** *An n-party protocol π, where a designated sender P_s holds an input $m \in \{0,1\}^*$, is a* t-secure *protocol for* broadcast with unanimous abort *(UA-BC, or UA-broadcast) if for every t-adversary the following hold except with negligible probability.*

- ***Termination.*** *Every honest party terminates.*
- ***Agreement.*** *All honest parties output the same value $y \in \{0,1\}^* \cup \{\perp\}$.*
- ***Validity.*** *If all parties are honest, then every party outputs m. If the sender is honest, then every honest party outputs m or $\perp$.*

A protocol for *broadcast* (BC) additionally guarantees that an honest sender makes every honest party output m, and a protocol for *broadcast with identifiable abort* (IA-BC, or IA-broadcast) lets the honest parties output $\perp$ only together with the identity of the same corrupted party [CCF+26], so both are UA-BC protocols.

If there is moreover an *a-priori* known bound R such that every honest party terminates within R rounds, except with negligible probability, we call π an *R-round* UA-BC protocol and R its *round complexity*. Such a protocol is said to have *deterministic termination*, and is sometimes called deterministic for short, although it may be randomized. A UA-BC protocol with *probabilistic*

termination [CCGZ21] has no such bound, and instead requires every honest party to terminate with probability 1. Its *expected round complexity* is the supremum, over all t-adversaries and all inputs, of the expected round in which the last honest party terminates. As a party outputs no later than it terminates, this is at least the expected output round of every honest party.

3 Lower Bound for Broadcast with Unanimous Abort

This section proves Theorem 1.1, following the outline of Section 1.2, which we briefly recall. Garay *et al.* [GKKO07] split the parties into sets on a line, with the sender's set at one end and any two adjacent sets jointly containing at least h parties, and consider the line execution, in which every message between non-adjacent sets is dropped. This execution is an attack on every pair of adjacent sets, so agreement gives all parties a common output, and validity at the sender's end makes it the sender's input. On the longest such line, however, the far party learns nothing about this input before round d. This line argument does not apply to UA-broadcast as it stands. Once a party is corrupted, validity allows the output $\perp$, so in the line execution all parties may output $\perp$, and the far party can do so without knowing the sender's input. The all-honest execution is the only one in which validity forces the output to be the sender's input, and our argument starts there.

It walks from the all-honest execution to the line execution through a chain of executions, each omitting a few more messages between non-adjacent sets than the one before. Every link of the chain, that is, every two consecutive executions, has a *witness*. This is a party at the end of the line opposite to the omitted messages, whose view is the same in both executions until the effect of the omissions reaches it. The effect moves one set per round and has about $n/(n-t)$ sets to cross, so this takes about $n/(n-t)$ rounds. If the witness outputs before that, it outputs the same value in both executions, so the common output stays the sender's input unless the witness is slow. At the end of the chain the far party would have to output the sender's input without knowing it, so some witness, or the far party itself, must be slow.

Section 3.1 defines the executions and the chain and proves the two lemmas that Section 4 uses as well. Section 3.2 proves the bound, with adversaries that find a slow party during the execution.

3.1 Executions on a Line

Lines, omission patterns and executions. Let P_s be the sender and $L \triangleq \lfloor (d+1)/2 \rfloor$. As $d+1 = \lceil n/h \rceil + \lfloor n/h \rfloor$ for $h \geq 2$, we have $L = \lfloor n/h \rfloor$ if $h \geq 2$ and $L = \lfloor n/2 \rfloor$ if $h = 1$. Our *main line* $(Q_1, \ldots, Q_{d+1})$ is a partition of $[n]$ into $d+1$ nonempty sets with $Q_1 = \{s\}$, so its ends are at distance d. For $h = 1$ the sets are singletons. For $h \geq 2$ we write $n = Lh + e$ with $0 \leq e < h$, and take L consecutive pairs of sets of sizes 1 and $h-1$, followed by a set of size e if $e > 0$. This gives $\lceil n/h \rceil + \lfloor n/h \rfloor = d+1$ sets, and any two adjacent sets together contain at least h parties, so an adversary may corrupt every party outside them. Figure 3 shows the main line for $h = 4$ and $n = 14$. In general we call a sequence $\Lambda = (H_1, \ldots, H_k)$ of nonempty sets that partition $[n]$ and satisfy $|H_a \cup H_{a+1}| \geq h$ for every $a \in [k-1]$ a *line*, and for $a \in [k]$ and $J \subseteq [k]$ we write $\mathrm{dist}(a, J) \triangleq \min_{c \in J} |a - c|$. Such a line contains $\lfloor k/2 \rfloor$ disjoint pairs of adjacent sets with at least h parties each, so $\lfloor k/2 \rfloor h \leq n$, and $\lfloor k/2 \rfloor h \leq n-1$ if k is odd, as the last set is then left over. This gives $k \leq \lceil n/h \rceil + \lfloor n/h \rfloor$, and as trivially $k \leq n$, every line has at most $d+1$ sets. Garay *et al.* [GKKO07] use $2n/h$ sets of size $h/2$, and our unequal sizes make the main line as long as possible. Section 4 uses lines other than the main one.

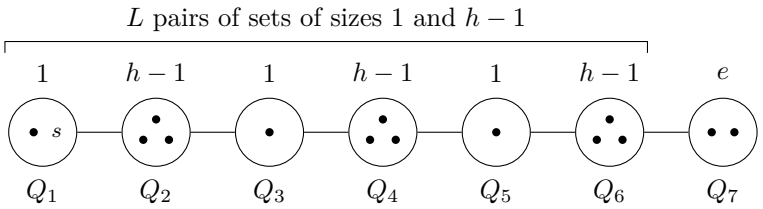


Figure 3: The main line for $h = 4$ and $n = 14$, so that $L = 3$, $e = 2$ and $d = 6$. Circles are sets and dots are parties. The sender alone forms Q_1, the L pairs of sets have sizes 1 and $h - 1$, and the last set holds the remaining e parties. Any two adjacent sets together contain at least h parties. It is also the line of Figure 6.

We index the message from P_u to P_v in round r by the triple (u, v, r), which we call a *slot*. We speak of slots rather than messages because the content of a message depends on the execution, whereas the sets of slots below are fixed before the execution. A slot (u, v, r) is a Λ-*cross slot* if P_u and P_v lie in sets of Λ that are neither equal nor adjacent, and simply a *cross slot* when Λ is the main line. An *omission pattern* is a set D of slots. An *execution* of π is the transcript of a run of π, namely the messages that every party sends and receives in every round, together with the output of every party. It is determined by the sender's input, the setup, the random tapes of the parties and the set of slots whose messages are not delivered. For fixed κ, n, t and π, we consider the random variables

$$\mathsf{Coins} \triangleq (B, \Sigma, U), \qquad U = (U_1, \ldots, U_n),$$

where B is the sender's input, a uniformly random bit, Σ is the setup of all parties, sampled from the setup distribution of the model, which may correlate the setups of different parties as a PKI does, and U_i is the random tape of P_i, a uniformly random string. The model samples B, Σ and $U_1, \ldots, U_n$ independently of each other, and our proofs use only that B is independent of (Σ, U). Unless stated otherwise, all probabilities and expectations are over Coins, in particular over a uniformly random input bit; a lower bound on an expected output round over Coins is a bound on the average over the two inputs, so it holds for at least one fixed input. Every omission pattern of this subsection and of Section 4 is fixed before Coins is sampled, and Section 3.2 says how it treats patterns chosen during the execution. For an omission pattern D, an input bit b, a setup σ and tapes $u = (u_1, \ldots, u_n)$, let $E_D(b, \sigma, u)$ be the execution on this input, setup and tapes in which all parties follow π, except that the messages of exactly the slots in D are not delivered, and let $E_D \triangleq E_D(\mathsf{Coins})$. So Coins is sampled once, every E_D is a deterministic function of the same Coins and of its fixed pattern D, and a different pattern does not trigger a fresh sampling. In particular, $E_\emptyset$ is the *all-honest execution*. For $\rho \geq 1$ let D_ρ be the set of cross slots of rounds $\rho, \rho + 1, \ldots$, so that $D_{\rho+1} \subseteq D_\rho$ and D_1 is the set of all cross slots. We call E_{D_1}, in which every set hears only from itself and its neighbors, the *line execution*. It is the execution that Garay *et al.* [GKKO07] analyze.

In these terms, an oblivious omission adversary (Section 2.1) fixes a set $I \subseteq [n]$ of corrupted parties and an omission pattern D, in which every slot has an endpoint in I, before Coins is sampled, and for each slot in D a corrupted endpoint withholds or ignores the message. The execution against it is E_D, whichever set I it corrupts. If $|I| \leq t$ and membership in D can be decided in polynomial time, as for every adversary we use, it is a t-adversary, so every guarantee of π holds against it.

One execution, many attacks. For a line $\Lambda = (H_1, \ldots, H_k)$, an omission pattern D of Λ-cross slots and $q \in [k - 1]$, let $\mathcal{A}^{\Lambda}_{D,q}$ be the oblivious omission adversary that leaves exactly the adjacent sets H_q and H_{q+1} honest and omits the slots in D. It is well defined since no Λ-cross slot has both

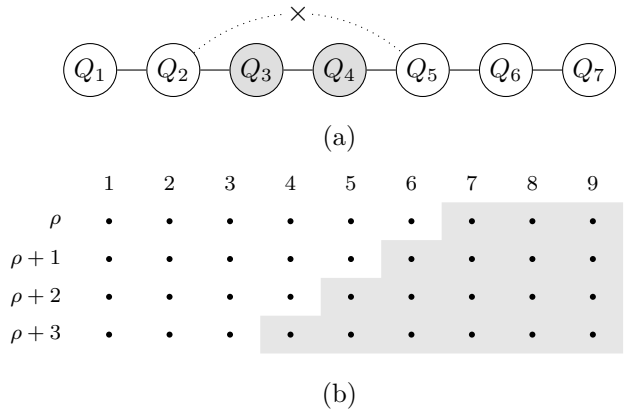


Figure 4: (a) A line with seven sets. Against $\mathcal{A}_{D,3}$ only Q_3 and Q_4 are honest, so every slot between non-adjacent sets, like the dotted one, has a corrupted endpoint and can be omitted. (b) Lemma 3.1 for a line with nine sets, a change in round ρ into the sets $J = \{7, 8, 9\}$, and all cross slots of the later rounds omitted. Rows are rounds and columns are sets. Views can differ only in the shaded region, which grows by one set per round, so a party in Q_1 keeps its view through round $\rho + \mathrm{dist}(1, J) - 1 = \rho + 5$.

endpoints in $H_q \cup H_{q+1}$, and it corrupts at most $n - h = t$ parties. The execution against $\mathcal{A}^{\Lambda}_{D,q}$ is E_D for every q, as it is for every oblivious omission adversary with the pattern D, whichever parties it corrupts. So the single execution E_D is at the same time the execution under many attacks, and a guarantee of π that holds against one of these attacks holds in E_D for the parties that are honest in that attack. This observation is at the core of our arguments. Garay *et al.* [GKKO07] use it for the line execution alone, whereas we use it for every execution of a chain and read the guarantees of all attacks on an execution at once, which is what lets us compare executions on the same randomness. On the event that agreement and termination hold against $\mathcal{A}^{\Lambda}_{D,q}$ for every $q \in [k-1]$, every party outputs in E_D and all outputs are equal, since the honest parties against $\mathcal{A}^{\Lambda}_{D,q}$ and against $\mathcal{A}^{\Lambda}_{D,q+1}$ share the set H_{q+1}. We write $Y(D)$ for this common output, a random variable that we use only on this event. This chaining of agreement along the line is from [GKKO07]. It says nothing about the value of $Y(D)$, which for UA-broadcast is the sender's input B if $D = \emptyset$, on the event that validity holds, and may be $\perp$ otherwise. For the main line we write $\mathcal{A}_{D,q}$ for $\mathcal{A}^{\Lambda}_{D,q}$. Figure 4(a) shows such an adversary.

Comparing executions. All executions E_D are functions of the same Coins, while the omission patterns are fixed in advance. So a comparison of two executions, such as that some party has the same view in both through some round, is a statement about the value of Coins, which holds or fails once this value is fixed. We first establish such comparisons for every value of Coins in an event of our choosing, and only then bound the probability of this event. The *view* of a party through round r consists of its input, its setup, its random tape and the messages it received in rounds 1 to r. It determines what the party sends in round $r+1$, and whether and what it outputs at the end of round r, so a message that is not delivered changes only the view of its receiver. We write $Y_i(D)$ for the output of P_i in E_D, with $Y_i(D) \triangleq \mathsf{none}$ if P_i never outputs, a symbol distinct from the abort output $\perp$ of the protocol, and $T_i(D)$ for its *output round* in E_D, the round at the end of which it outputs, with $T_i(D) = 0$ if P_i outputs before the first round and $T_i(D) = \infty$ if it never outputs. Like E_D, both are random variables, that is, functions of Coins. Our convention is that random variables are uppercase and their values lowercase. When we fix a value $\mathsf{coins} = (b, \sigma, u)$

of Coins, every E_D is the fixed execution $E_D(\mathsf{coins})$, and b, $y_i(D)$ and $y(D)$ denote the values of B, $Y_i(D)$ and $Y(D)$, while a condition on an output round is an event, such as $T_i(D) \le L$, that holds or fails at coins. Fixing a value is not conditioning on an event, and all probabilities stay over Coins.

Propagation of differences along a line. Consider two executions on the same setup and random tapes that differ in the sender's input or in the omission pattern, and suppose that from some round on both omit all Λ-cross slots. From that round on every party receives messages only from its own and the neighboring sets, so a difference between the two executions spreads along Λ by at most one set per round, as in Figure 4(b). Without this omission a difference could jump along the line in one round.

Lemma 3.1. *Let $\Lambda = (H_1, \ldots, H_k)$ be a line and $\rho \le \rho'$, where ρ' may be ∞, and consider two executions of π with the same setup and random tapes, possibly on different sender inputs, that omit the same slots in rounds $\rho+1, \ldots, \rho'$, including all Λ-cross slots. If through round ρ only parties in the sets H_c with $c \in J$ have different views, then a party in a set H_a with $a \notin J$ has the same view in both executions through round* $\min\{\rho', \rho + \mathrm{dist}(a, J) - 1\}$.

Proof. We show by induction on $r \in \{\rho, \ldots, \rho'\}$ that every party in a set H_a with $\mathrm{dist}(a, J) > r - \rho$ has the same view through round r in both executions. For $r = \rho$ this is the hypothesis, since $\mathrm{dist}(a, J) > 0$ means $a \notin J$. Let $r < \rho'$ and let $\mathrm{dist}(a, J) > r + 1 - \rho$. In round $r+1$ a party in H_a receives messages only from $H_{a-1} \cup H_a \cup H_{a+1}$, since both executions omit all Λ-cross slots of this round, and these three sets are at distance greater than $r - \rho$ from J. By the induction hypothesis their parties have the same views through round r in both executions, so they send the same messages in round $r+1$. As both executions omit the same slots of round $r+1$, the party in H_a receives the same messages in round $r+1$ and has the same view through round $r+1$. For $a \notin J$ this gives the same view through every round $r \le \rho'$ with $r - \rho < \mathrm{dist}(a, J)$, which is the claim. □

We use the lemma in two ways. First, let D and D' be omission patterns that differ only in slots of round $\rho \ge 1$ with receivers in $\bigcup_{c \in J} H_c$ and that both contain all Λ-cross slots of the rounds $\rho+1, \ldots, \rho'$. Then, for every value of Coins, E_D and $E_{D'}$ deliver the same messages before round ρ, so every party has the same view through round $\rho - 1$ in both and sends the same messages in round ρ, and in round ρ the two executions differ only in messages into $\bigcup_{c \in J} H_c$. So the lemma applies, and a party in H_a with $a \notin J$ has the same view in E_D and in $E_{D'}$ through round $\min\{\rho', \rho + \mathrm{dist}(a, J) - 1\}$. This is how the links of our chains are built. Second, in the line execution the sender's input initially affects only the sender. For a fixed setup σ and tapes u, apply the lemma to the main line and to $E_{D_1}(0, \sigma, u)$ and $E_{D_1}(1, \sigma, u)$, with $\rho = 0$, $J = \{1\}$ and $\rho' = \infty$. A party P_i in Q_j with $j \ge 2$ has the same view through round $j-2$ in both, so its view through round $j-2$ in E_{D_1} is a function of (Σ, U) alone and hence independent of B. Let $r \le j-2$. Whether $T_i(D_1) \le r$, and the value of $Y_i(D_1)$ if so, are functions of this view. As B is uniform and independent of the view, for every value v of the view with $T_i(D_1) \le r$ the output is a value $f(v)$ with $\Pr[B = f(v)] \le \frac{1}{2}$, and summing over v gives

$$\Pr[Y_i(D_1) = B \wedge T_i(D_1) \le r] \;\le\; \tfrac{1}{2}\Pr[T_i(D_1) \le r] \qquad \text{for } i \in Q_j,\ j \ge 2 \text{ and } r \le j-2. \tag{1}$$

This is the blindness of the far party in [GKKO07], and $j = d+1$ gives it through round $d-1$, which yields the bound of d rounds for broadcast.

A uniform negligible-error bound. Our lower bounds compare polynomially many executions on the same Coins, and we need the properties of π to hold in all of them at once. Each $\mathcal{A}^{\Lambda}_{D,q}$ is a t-adversary, but the negligible error bound may depend on the adversary, and the list of adversaries grows with κ. The next lemma gives one bound for the whole list.

Lemma 3.2. *If a property of π holds except with negligible probability against every oblivious omission t-adversary, then, except with negligible probability over* Coins*, it holds at once against all adversaries of any polynomial-time computable list of polynomially many such adversaries.*

Proof. Let $\mathcal{A}_1, \ldots, \mathcal{A}_N$ be such a list. Let an adversary pick $i \in [N]$ uniformly at random with its own coins, independently of Coins, and then act as $\mathcal{A}_i$. This is again an oblivious omission t-adversary, as it makes its choice before Coins is sampled, and conditioned on its choice, the execution is that of $\mathcal{A}_i$. So if the property fails against $\mathcal{A}_i$ with probability ε_i, it fails against the mixture with probability $\frac{1}{N}\sum_i \varepsilon_i$, which is negligible. By a union bound, it fails against some $\mathcal{A}_i$ with probability at most $\sum_i \varepsilon_i$, which is N times a negligible quantity and hence negligible. □

The chain on the main line. We call $Q_1, \ldots, Q_L$ the *left half* of the main line and $Q_{L+1}, \ldots, Q_{d+1}$ its *right half*. The sender's set is at distance $\mathrm{dist}(1, \{L+1, \ldots, d+1\}) = L$ from the right half, and the last set at distance $\mathrm{dist}(d+1, [L]) = d+1-L \geq L$ from the left half, as $d+1 \geq 2L$. For $\rho \in [L]$, let D'_ρ be the union of $D_{\rho+1}$ and the set of cross slots of round ρ whose receiver is in the right half, so that $D_{\rho+1} \subseteq D'_\rho \subseteq D_\rho$. We call two consecutive executions of a chain a *link*. For $k \in [L]$, the links of the chain

$$E_{D_{k+1}} \longrightarrow E_{D'_k} \longrightarrow E_{D_k} \longrightarrow \cdots \longrightarrow E_{D_2} \longrightarrow E_{D'_1} \longrightarrow E_{D_1} \tag{2}$$

omit the cross slots of the rounds $k, k-1, \ldots, 1$, one round at a time, and each round in two halves, first into the right half and then into the left half. The rounds are cut backward because Lemma 3.1 needs all cross slots of the later rounds omitted in both executions of a link, and a round is cut in halves because omitting a whole round changes views next to every set, whereas each half leaves the opposite end of the line at distance at least L. Figure 5 shows the slots that E_{D_3}, $E_{D'_2}$ and E_{D_2} deliver, and Figure 6 shows the chain for seven sets and $k = 2$.

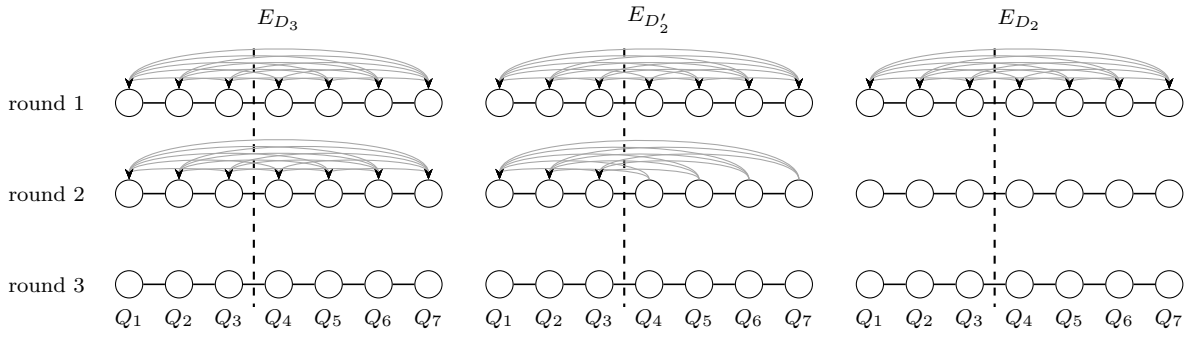


Figure 5: The slots delivered in E_{D_3}, $E_{D'_2}$ and E_{D_2}, three consecutive executions of the chain, on the main line of Figure 3 with $L = 3$. Each row is a round, each circle a set of the line, and the dashed line separates the left half Q_1, Q_2, Q_3 from the right half $Q_4, \ldots, Q_7$. A segment between adjacent sets stands for the slots between their parties in that round in both directions, and an arc between non-adjacent sets for those in the directions of its arrowheads. Slots within a set are always delivered and are not drawn, and a missing arc is an omitted cross slot. D_3 omits every cross slot of rounds $3, 4, \ldots$, so rounds 1 and 2 deliver everything and from round 3 on only adjacent sets communicate. D'_2 also omits the cross slots of round 2 whose receiver is in the right half, so in round 2 the right half still sends to the left half while the left half no longer reaches the right half. D_2 omits every cross slot of rounds $2, 3, \ldots$. Later rounds look like round 3.

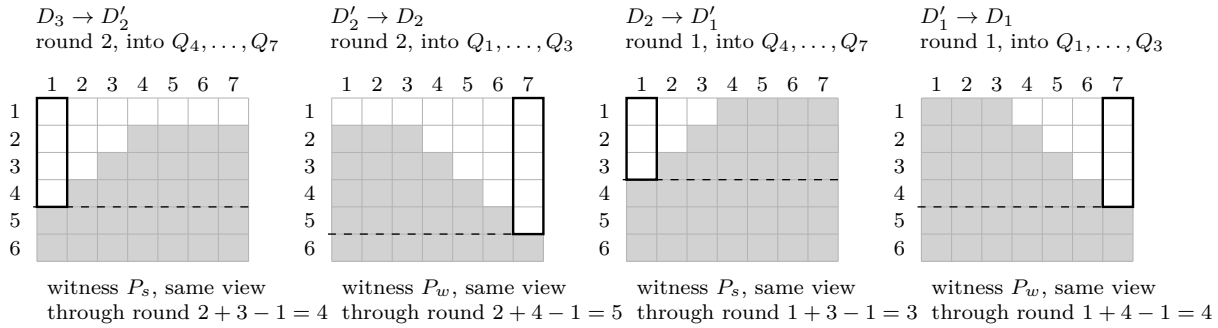


Figure 6: The chain (2) for $d = 6$, that is, seven sets and $L = 3$, with $k = 2$. Its four links follow from left to right. Each panel names the cross slots that the link omits in addition, and shows a grid with rounds as rows and sets as columns, in which a shaded cell means that a party in that set may have a different view at the end of that round in the two executions. The shaded region starts in round ρ in the changed half and grows by one set per round, since all cross slots of the later rounds are omitted in both executions. The witness, boxed, is at the far end of the line and keeps its view through the round of the dashed line. Every horizon is at least $L = 3$, and the sender's horizon for $\rho = 1$ equals L.

3.2 The Lower Bound

The adversaries of Section 3.1 fix their omission pattern before Coins is sampled. Along the chain (2) some witness, or the far party at its end, must be slow, but which one may depend on Coins, so an adversary that fixes its omissions in advance may rarely meet a slow party. Against such adversaries the chain gives only $\Omega(\sqrt{n/(n-t)})$ rounds (Appendix A). This subsection therefore lets the adversary choose during the execution the round from which it omits all cross slots of the main line, which we call *cutting* at that round, and shows that UA-broadcast then takes at least $t/(2(n-t))$ rounds in expectation. The corrupted parties still follow the protocol and still omit only cross slots, so the adversary realizes one of the executions E_{D_ρ} of Section 3.1 or the all-honest execution.

The proof has three steps, as outlined in Section 1.2. First, it considers a party P_v next to the sender and a party P_w at the far end, and shows by agreement that one of them is slow when the line is cut at a suitable round, or not at all, where the round depends on Coins. Second, it defines two adversaries, one keeping P_w honest and one keeping P_v honest, which recognize this round during the execution and cut there (Claim 3.3). This is the one place where rushing is used. Third, it adds up the probabilities of the two attacks.

Rushing omission adversaries. This subsection uses rushing omission adversaries (Section 2.1) that are deterministic. Such an adversary fixes the set of corrupted parties before Coins is sampled, but chooses the omissions of a round as a function of what the corrupted parties know by then, that is, of B if the sender is corrupted, of the setup and random tapes of the corrupted parties, and of all messages they have received, including those of that round. An omission adversary that runs in polynomial time and corrupts at most t parties is a t-adversary. For a value $\mathsf{coins} = (b, \sigma, u)$ of Coins, let $D(\mathsf{coins})$ be the set of all slots whose messages the adversary withholds or ignores in the execution on coins. Section 3.1 defines $E_D(b, \sigma, u)$ for every omission pattern D, while every pattern it considers is fixed in advance, and the execution against the adversary on coins is $E_{D(\mathsf{coins})}(b, \sigma, u)$. Indeed, in both executions all parties follow the protocol and exactly the messages of the slots in $D(\mathsf{coins})$ are not delivered, so an induction on the rounds shows that the two executions deliver the same messages in every round. Hence such an adversary realizes, on every value of Coins, one execution of the family of Section 3.1, and the pattern it realizes is a function of Coins. The adversaries below omit, from some round on, exactly the cross slots of the main line, so they realize E_{D_ρ} with a round ρ chosen during the execution, or $E_\emptyset$ if they never cut.

Theorem 1.1. *Let $t \le n-2$ and let π be any t-secure UA-BC protocol against rushing omission adversaries, then its expected round complexity is at least $\frac{t}{2(n-t)}$.*

Proof. We use the probability space and notation of Section 3.1, with $L = \lfloor (d+1)/2 \rfloor$, the left and right halves of the main line and the patterns D'_ρ. If $L = 1$, then $n/(n-t) < 2$ as $t \le n-2$, so the bound is below $1/2$, and the all-honest execution gives it, since a party other than the sender outputs B in $E_\emptyset$ except with negligible probability, while its view before the first round is independent of B, so its expected output round is at least $1 - \mathsf{negl}(\kappa)$. So let $L \ge 2$. Then $d \ge 3$, and $d \ge L+1$ as $d+1 \ge 2L$. We show that against one of two polynomial-time rushing omission adversaries some honest party has not output by round $L-1$ with probability at least $(1-2\delta)/2$, where δ is negligible, so that its expected output round is at least $L(1-2\delta)/2$. This gives the theorem. Indeed, for $t \le n-2$ we have $L = \lfloor n/(n-t) \rfloor \ge (t+1)/(n-t)$, so $L(1-2\delta)/2 \ge t/(2(n-t))$ once the loss δL is at most $1/(2(n-t))$. As δ is negligible and $L(n-t) \le n$, this holds for all sufficiently large κ. We fix $v \in Q_2$ and $w \in Q_{d+1}$. We use P_v rather than the sender as the witness on the left so that the sender is corrupted in both attacks below, whose adversaries need to know B. Figure 7 shows the two parties and the honest sets of the two attacks.

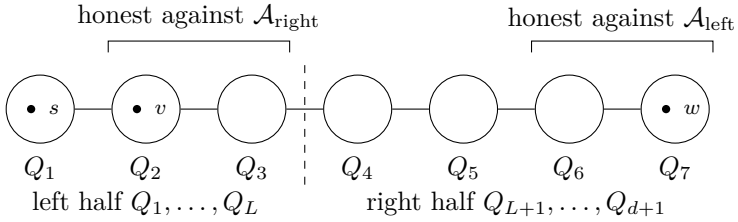


Figure 7: The parties and attacks of the proof of Theorem 1.1, on the main line with seven sets, so that $L = 3$ and $d = 6$, as in Figure 6. Only the named parties are drawn. The sender P_s forms Q_1, P_v lies in Q_2, next to the sender, and P_w lies in Q_{d+1}, at the far end. The adversary $\mathcal{A}_{\text{left}}$ keeps $Q_d \cup Q_{d+1}$ honest and $\mathcal{A}_{\text{right}}$ keeps $Q_2 \cup Q_3$ honest, so P_w is honest against the first, P_v against the second, and the sender is corrupted in both. Each of them cuts the line, that is, omits all cross slots from some round on, at a round that it chooses during the execution from what the party at the other end would output. The dashed line separates the left half from the right half of the line.

Statuses. For a pattern D and a party P_i, the *status* of P_i in E_D is yes if $T_i(D) \le L-1$ and $Y_i(D) = B$, it is no if $T_i(D) \le L-1$ and $Y_i(D) \ne B$, and it is slow if $T_i(D) > L-1$. It is a

random variable with values in $\{\mathsf{yes}, \mathsf{no}, \mathsf{slow}\}$, since $T_i(D)$, $Y_i(D)$ and B are random variables, and it is a function of B and of the view of P_i through round $L-1$ in E_D. As in Section 3.1, when we fix a value coins of Coins the status is its value at coins, and a condition on statuses, such as $Z_k = \mathsf{yes}$ below, is an event that holds or fails at coins. Let $\mathcal{D} \triangleq \{D_1, \ldots, D_L, D'_1, \ldots, D'_{L-1}\}$. We consider the $N \triangleq (2L-1)d+1 \leq n^2$ adversaries $\mathcal{A}_{D,q}$ with $D \in \mathcal{D}$ and $q \in [d]$, together with the adversary that corrupts no party, and let Good be the event that agreement and termination hold against every $\mathcal{A}_{D,q}$ with $D \in \mathcal{D}$, that validity holds against $\mathcal{A}_{D_1,1}$, and that validity holds in $E_\emptyset$. By Lemma 3.2, $\Pr\left[\overline{\mathsf{Good}}\right] \leq \delta$ for a negligible δ. On Good, all parties have the common output $Y(D)$ in E_D for every $D \in \mathcal{D}$, by Section 3.1, so no execution E_D with $D \in \mathcal{D}$ has one party with status yes and another with status no.

The next two paragraphs do not involve the attacks yet. They compare the statuses of P_w and P_v across the fixed executions E_D with $D \in \mathcal{D}$ and show that, except with negligible probability, one of the two parties is slow in an execution E_{D_ρ} or in $E_\emptyset$, where the round ρ of the cut depends on Coins. The two adversaries of the proof, defined afterwards, realize exactly this execution.

The chain of statuses. Consider the sequence $Z_1, \ldots, Z_{2L-1}$, where $Z_{2\rho-1}$ is the status of P_w in E_{D_ρ} for $\rho \in [L]$ and $Z_{2\rho}$ is the status of P_v in $E_{D_{\rho+1}}$ for $\rho \in [L-1]$. We show that two consecutive entries are the statuses of P_w and P_v in one execution E_D with $D \in \mathcal{D}$. Let $\rho \in [L-1]$. The patterns $D_{\rho+1}$ and D'_ρ differ only in cross slots of round ρ with receivers in the right half and both contain all cross slots of the later rounds, so Lemma 3.1, in its first use in Section 3.1 with $\rho' = \infty$ and $J = \{L+1, \ldots, d+1\}$, shows that P_v has the same view through round $\rho + \mathrm{dist}(2, J) - 1 = \rho + L - 2 \geq L-1$ in $E_{D_{\rho+1}}$ and in $E_{D'_\rho}$. So it either outputs the same value by round $L-1$ in both executions or outputs after round $L-1$ in both, and in either case it has the same status in both. Likewise D'_ρ and D_ρ differ only in cross slots of round ρ with receivers in the left half, and the lemma with $J = [L]$ shows that P_w has the same view through round $\rho + d - L \geq L-1$ in $E_{D'_\rho}$ and in E_{D_ρ}, as $d+1 \geq 2L$. So $Z_{2\rho-1}$ and $Z_{2\rho}$ are the statuses of P_w and P_v in $E_{D'_\rho}$, and $Z_{2\rho}$ and $Z_{2\rho+1}$ are their statuses in $E_{D_{\rho+1}}$. These are the links of the chain (2) for $k = L-1$, read for statuses and with P_v in the place of the sender. As P_v is one set closer to the right half than the sender, it keeps its view one round less, through round $\rho + L - 2$ rather than $\rho + L - 1$, which is why the statuses use the deadline $L-1$ in place of L.

We now show that the first entry of the sequence $Z_1, \ldots, Z_{2L-1}$ is not yes, except with negligible probability, and that its last entry is yes or slow, the former by blindness and the latter by validity. The first entry is the status of P_w in the line execution. As shown in Section 3.1 before the bound (1), the view of P_w through round $d-1$ in E_{D_1} is a function of (Σ, U) alone and hence independent of B, and $L-1 \leq d-1$. We bound $\Pr[Z_1 = \mathsf{yes}]$ by showing that P_w outputs a bit equal to B by round $L-1$ exactly as often as it outputs a bit different from B by round $L-1$, since its view through round $L-1$ is independent of B, and that the latter has probability at most δ, since validity leaves P_w only the outputs B and $\perp$. Let V be the view of P_w through round $L-1$ in E_{D_1}. Whether $T_w(D_1) \leq L-1$, and the value of $Y_w(D_1)$ if so, are functions of V. Let X be the set of values of V on which P_w outputs a bit by round $L-1$, and for $x \in X$ let $f(x)$ be this bit. Then $Z_1 = \mathsf{yes}$ is the event $V \in X \wedge B = f(V)$, and P_w outputs a bit different from B by round $L-1$ exactly on the event $V \in X \wedge B \neq f(V)$. As B is a uniform bit independent of V, we have $\Pr[V = x \wedge B = f(x)] = \Pr[V = x \wedge B \neq f(x)] = \frac{1}{2}\Pr[V = x]$ for every $x \in X$, and summing over $x \in X$ shows that both events have probability $\frac{1}{2}\Pr[V \in X]$. The second event is contained in the event $Y_w(D_1) \notin \{B, \perp\}$, which has probability at most δ, since on Good every party outputs the common output $Y(D_1)$ in E_{D_1}, and the sender and the parties in Q_2 are honest against $\mathcal{A}_{D_1,1}$, so validity gives $Y(D_1) \in \{B, \perp\}$. Together,

$$\Pr[Z_1 = \mathsf{yes}] \;=\; \Pr[T_w(D_1) \leq L-1 \wedge Y_w(D_1) \in \{0,1\} \setminus \{B\}] \;\leq\; \Pr[Y_w(D_1) \notin \{B, \perp\}] \;\leq\; \delta.$$

The bound (1) alone gives $\Pr[Z_1 = \mathsf{yes}] \le \frac{1}{2}$, which would halve the constant of the theorem. The last entry is the status of P_w in E_{D_L}, which delivers the same messages as $E_\emptyset$ in the first $L-1$ rounds, so it is the status of P_w in $E_\emptyset$. On Good, validity gives $Y_w(\emptyset) = B$, so Z_{2L-1} is yes or slow.

The slow party before the first yes**.** Let R_{left} be the least $\rho \in [L-1]$ such that P_v has status yes in $E_{D_{\rho+1}}$, and let R_{right} be the least $\rho \in \{2, \ldots, L\}$ such that P_w has status yes in E_{D_ρ}, each ∞ if no such ρ exists. Both are random variables, and the two adversaries below cut at these rounds. We set $D_\infty \triangleq \emptyset$ and write $T_w(D_{R_{\text{left}}})$ for $T_w(D_\rho)$ on the event $R_{\text{left}} = \rho$, for $\rho \in [L-1] \cup \{\infty\}$, and likewise $T_v(D_{R_{\text{right}}})$. We show that on $\mathsf{Good} \cap \{Z_1 \neq \mathsf{yes}\}$, an event of probability at least $1 - 2\delta$, either P_w has not output by round $L-1$ in $E_{D_{R_{\text{left}}}}$, or $R_{\text{right}} < \infty$ and P_v has not output by round $L-1$ in $E_{D_{R_{\text{right}}}}$, that is,

$$\mathsf{Good} \cap \{Z_1 \neq \mathsf{yes}\} \;\subseteq\; \{T_w(D_{R_{\text{left}}}) > L-1\} \;\cup\; \{R_{\text{right}} < \infty \wedge T_v(D_{R_{\text{right}}}) > L-1\}. \tag{3}$$

Fix a value coins of Coins in Good with $Z_1 \neq \mathsf{yes}$, and look at the first yes in $Z_1, \ldots, Z_{2L-1}$. There are three cases. If there is no yes, then $R_{\text{left}} = \infty$, and Z_{2L-1}, the status of P_w in $E_\emptyset$, is slow, as it is yes or slow by validity. So $T_w(\emptyset) > L-1$, which is the first event of (3) as $D_\infty = \emptyset$. If the first yes is an even entry $Z_{2\rho}$, the status of P_v in $E_{D_{\rho+1}}$, then $R_{\text{left}} = \rho$. The entry before it, $Z_{2\rho-1}$, is the status of P_w in E_{D_ρ}, and as shown above the two entries are the statuses of P_w and P_v in the one execution $E_{D'_\rho}$. On Good no execution E_D with $D \in \mathcal{D}$ has one party with status yes and another with status no, so $Z_{2\rho-1}$ is not no, and it is not yes as $Z_{2\rho}$ is the first yes. So $Z_{2\rho-1} = \mathsf{slow}$, that is, $T_w(D_\rho) > L-1$, which is again the first event of (3). If the first yes is an odd entry $Z_{2\rho+1}$, the status of P_w in $E_{D_{\rho+1}}$, then $\rho \ge 1$ as $Z_1 \neq \mathsf{yes}$, and $R_{\text{right}} = \rho + 1$. The entry before it, $Z_{2\rho}$, is the status of P_v in the same execution $E_{D_{\rho+1}}$, so by the same argument $Z_{2\rho} = \mathsf{slow}$, that is, $T_v(D_{\rho+1}) > L-1$, which is the second event of (3).

So far we have shown that, except with negligible probability, P_w has not output by round $L-1$ in $E_{D_{R_{\text{left}}}}$ or P_v has not output by round $L-1$ in $E_{D_{R_{\text{right}}}}$, where R_{left} and R_{right} depend on Coins. We now define two adversaries that realize these two executions, by cutting the line at round R_{left} with P_w honest and at round R_{right} with P_v honest. The point is that they can determine these rounds during the execution, which Claim 3.3 shows, and the accounting at the end then turns (3) into the bound.

The two adversaries. Let $\mathcal{A}_{\text{left}}$ corrupt every party outside $Q_d \cup Q_{d+1}$ and let $\mathcal{A}_{\text{right}}$ corrupt every party outside $Q_2 \cup Q_3$. Both corrupt at most $n - h = t$ parties, as adjacent sets of the main line together contain at least h parties, and both corrupt the sender, so both know B. In every round $\rho \in [L-1]$, as long as it has not cut, $\mathcal{A}_{\text{left}}$ decides whether $R_{\text{left}} = \rho$, and if so it cuts at round ρ, that is, it selects all cross slots of round ρ and of every later round, and otherwise it selects nothing. Likewise $\mathcal{A}_{\text{right}}$ decides in every round $\rho \in \{2, \ldots, L\}$ whether $R_{\text{right}} = \rho$, and cuts at round R_{right}. Claim 3.3 below shows that the two adversaries can make these decisions, so on every value of Coins, $\mathcal{A}_{\text{left}}$ realizes $E_{D_{R_{\text{left}}}}$ with P_w honest and $\mathcal{A}_{\text{right}}$ realizes $E_{D_{R_{\text{right}}}}$ with P_v honest, where $D_\infty = \emptyset$ is the pattern of an adversary that never cuts.

Claim 3.3. *The statuses that the two adversaries test can be computed as follows.*

(1) For $\rho \in [L-1]$, the status of P_v in $E_{D_{\rho+1}}$ is a polynomial-time computable function of B, of the setup and random tapes of the parties outside $Q_d \cup Q_{d+1}$, and of the messages that the parties in $Q_d \cup Q_{d+1}$ send to parties outside $Q_d \cup Q_{d+1}$ in rounds 1 to ρ of $E_\emptyset$.

(2) For $\rho \in \{2, \ldots, L\}$, the status of P_w in E_{D_ρ} is a polynomial-time computable function of B, of the setup and random tapes of the parties outside $Q_2 \cup Q_3$, and of the messages that the parties in $Q_2 \cup Q_3$ send to parties outside $Q_2 \cup Q_3$ in rounds 1 to $\rho - 1$ of $E_\emptyset$.

Proof. We prove part (1). In $E_{D_{\rho+1}}$ the parties in $Q_d \cup Q_{d+1}$ still send messages after round ρ, to their neighbor Q_{d-1}, so the view of P_v through round $L-1$ is not obviously determined by the data of the claim. We therefore pass to an execution $E_{D''}$ in which nothing that these parties send after round ρ is delivered outside $Q_d \cup Q_{d+1}$, where this is clear, and then show that P_v has the same view through round $L-1$ in $E_{D''}$ and in $E_{D_{\rho+1}}$, since the omitted messages into Q_{d-1} cannot affect P_v by round $L-1$. Let D'' be the union of $D_{\rho+1}$ and the set of slots of rounds $\rho+1, \ldots, L-1$ from parties in $Q_d \cup Q_{d+1}$ to parties outside $Q_d \cup Q_{d+1}$. The patterns $\emptyset$ and D'' agree on rounds 1 to ρ, so $E_{D''}$ and $E_\emptyset$ deliver the same messages in these rounds, every party has the same view through round ρ in both, and the parties in $Q_d \cup Q_{d+1}$ send the same messages in rounds 1 to ρ in both. In $E_{D''}$ no message that a party in $Q_d \cup Q_{d+1}$ sends in a round in $\{\rho+1, \ldots, L-1\}$ is delivered outside $Q_d \cup Q_{d+1}$. So an induction on the rounds shows that the view through round $L-1$ of every party outside $Q_d \cup Q_{d+1}$ in $E_{D''}$ is a function of the setup and random tapes of these parties, of B, and of the messages named in the claim, and that running π for these parties on these data computes it in polynomial time. This holds in particular for the status of P_v in $E_{D''}$. It remains to show that P_v has the same view through round $L-1$ in $E_{D''}$ and in $E_{D_{\rho+1}}$. If $\rho+1 > L-1$ the two patterns are equal. Otherwise $\rho \leq L-2$, so $L \geq 3$ and $d \geq 2L-1 \geq 5$. For $r \in \{\rho, \ldots, L-1\}$ let $D^{(r)}$ be the union of $D_{\rho+1}$ and the slots of D'' of rounds $\rho+1$ to r, so that $D^{(\rho)} = D_{\rho+1}$ and $D^{(L-1)} = D''$. For $r \in \{\rho+1, \ldots, L-1\}$, the patterns $D^{(r-1)}$ and $D^{(r)}$ differ only in slots of round r with receivers in Q_{d-1}, since a slot of D'' of round r with its receiver elsewhere outside $Q_d \cup Q_{d+1}$ is a cross slot and lies in $D_{\rho+1}$, and both contain all cross slots of the rounds after r. So Lemma 3.1, in its first use in Section 3.1 with $\rho' = \infty$ and $J = \{d-1\}$, shows that P_v has the same view in $E_{D^{(r-1)}}$ and in $E_{D^{(r)}}$ through round $r + \mathrm{dist}(2, \{d-1\}) - 1 = r + d - 4 \geq \rho + d - 3 \geq d - 2 \geq L - 1$, as $\rho \geq 1$ and $d \geq L+1$. Composing over r gives part (1).

Part (2) is proved in the same way with D'' the union of D_ρ and the slots of rounds $\rho, \ldots, L-1$ from parties in $Q_2 \cup Q_3$ to parties outside $Q_2 \cup Q_3$, with rounds 1 to $\rho-1$ in the place of rounds 1 to ρ, and with the receivers of the changed slots in $Q_1 \cup Q_4$. If $\rho > L-1$ the patterns are equal, and otherwise $L \geq 3$, $d \geq 5$ and $\mathrm{dist}(d+1, \{1,4\}) = d-3$, so P_w keeps its view through round $r + d - 4 \geq \rho + d - 4 \geq d - 2 \geq L - 1$, as $\rho \geq 2$. □

Against $\mathcal{A}_{\text{left}}$, as long as it has not cut, the parties in $Q_d \cup Q_{d+1}$ send in rounds 1 to ρ the messages they send in $E_\emptyset$. Part (1) of Claim 3.3 needs those of these messages that are sent to corrupted parties, including the messages of round ρ itself, and $\mathcal{A}_{\text{left}}$ selects the slots of round ρ only after the honest parties have sent them, since it is rushing. This is the one place where the rushing capability is used. With these messages, B and the setup and random tapes of the corrupted parties, $\mathcal{A}_{\text{left}}$ computes the status of P_v in $E_{D_{\rho+1}}$ by the claim, and it cuts if the status is yes, so it cuts exactly at R_{left}. Against $\mathcal{A}_{\text{right}}$, the same holds at the start of round $\rho \geq 2$ for the messages of rounds 1 to $\rho-1$, which part (2) needs and which have all been sent before round ρ starts, so $\mathcal{A}_{\text{right}}$ computes the status of P_w in E_{D_ρ} and cuts exactly at R_{right} without the rushing capability. The difference is that $\mathcal{A}_{\text{left}}$ cuts at round ρ on the strength of a status under the cut at round $\rho+1$, which depends on one more round of honest messages, while $\mathcal{A}_{\text{right}}$ cuts at round ρ on a status under the cut at round ρ itself. Both adversaries run in polynomial time, by Claim 3.3 and since the corrupted parties follow π, so both are rushing omission t-adversaries. Section 2 assumes that honest parties run in polynomial time in total. A protocol with probabilistic termination may instead be polynomial-time in each round and run for an unbounded number of rounds, and then our adversaries, whose corrupted parties follow the protocol, are not polynomial-time. In this case let $P \triangleq n^3\kappa^\alpha$ for a constant α. The corrupted parties of every adversary of this proof, $\mathcal{A}_{\text{left}}$ and $\mathcal{A}_{\text{right}}$ included, stop after round P, which makes all of them polynomial-time, as the statuses they compute concern rounds before P, and which changes no view through round P. With C

the largest expected output round of a party honest against one of the stopped adversaries, Good also requires every party to output by round P in every E_D with $D \in \mathcal{D}$, which adds at most $n(2L-1)C/(P+1)$ to δ by Markov's inequality, and the bound $L(1-2\delta)/2$ then holds for C up to an additive $O(L\kappa^{-\alpha})$. For a large enough constant α, this term and the loss δL together stay below $1/(2(n-t))$ for all sufficiently large κ, so the theorem holds in this case as well.

The accounting. Let p_{left} be the probability that P_w has not output by round $L-1$ against $\mathcal{A}_{\text{left}}$ and p_{right} the probability that P_v has not output by round $L-1$ against $\mathcal{A}_{\text{right}}$. Since $\mathcal{A}_{\text{left}}$ realizes $E_{D_{R_{\text{left}}}}$ and $\mathcal{A}_{\text{right}}$ realizes $E_{D_{R_{\text{right}}}}$ on every value of Coins, the first event on the right of (3) has probability at most p_{left} and the second at most p_{right}. The event on the left of (3) is the complement of $\{Z_1 = \mathsf{yes}\} \cup \overline{\mathsf{Good}}$, and each of these two events has probability at most δ, so it has probability at least $1-2\delta$ and (3) gives

$$1 - 2\delta \;\leq\; p_{\text{left}} + p_{\text{right}}.$$

So the larger of p_{left} and p_{right} is at least $(1-2\delta)/2$. That is, P_w against $\mathcal{A}_{\text{left}}$ or P_v against $\mathcal{A}_{\text{right}}$, each honest in its attack, has not output by round $L-1$ with probability at least $(1-2\delta)/2$, and its expected output round is then at least $L(1-2\delta)/2$. As all probabilities are over a uniformly random input bit, the same bound holds on at least one fixed input, as noted in Section 3.1. □

The guarantees of π enter the proof only through Good, that is, agreement and termination against the oblivious omission adversaries $\mathcal{A}_{D,q}$ with $D \in \mathcal{D}$, validity against $\mathcal{A}_{D_1,1}$ and validity in the all-honest execution. The proof assumes nothing about π against $\mathcal{A}_{\text{left}}$ and $\mathcal{A}_{\text{right}}$. Their only role is to realize, on each value of Coins, one of the executions E_{D_ρ} or $E_\emptyset$ that Good speaks about, with the slow party honest. Validity against $\mathcal{A}_{D_1,1}$, that is, the honest-sender clause of Definition 2.1 in the line execution, serves only the constant. With the bound (1) in its place the proof uses validity in the all-honest execution alone, as the proof of Section 4 does, and gives $L(1-2\delta)/4$. If the sender's input is instead uniform over $\{0,1\}^\kappa$, which Definition 2.1 allows, the derivation of the bound (1) in Section 3.1 gives it with $2^{-\kappa}$ in place of $\frac{1}{2}$, and the proof gives $L(1-\delta-2^{-\kappa})/2$ while using validity in the all-honest execution alone.

Tightness. The protocol of Dolev and Strong [DS83] terminates deterministically in $t+1$ rounds in the signature-PKI model. For every constant number $h = n - t$ of honest parties it takes $\Theta(n) = \Theta(t/(n-t))$ rounds, so the bound of Theorem 1.1 is asymptotically tight for $h = O(1)$. It is a broadcast protocol secure against Byzantine adversaries, which is more than the lower bound asks for in either respect, so the comparison bounds the round complexity of UA-broadcast from both sides.

4 Lower Bound for Deterministic Broadcast with Unanimous Abort

This section proves Theorem 1.2, following the outline of Section 1.2. With deterministic termination, every party outputs by round R in each execution we consider, except with negligible probability. A link then keeps the common output whenever its witness, the party far from the omitted messages, has the same view through round R in both executions, and it loses only the negligible probability with which the guarantees fail. Unlike in Section 3, where a link fails when its witness is slow, the chain may therefore be polynomially long.

By Lemma 3.1, the witness keeps its view through round R if the link changes messages of round r only into a set at distance at least $R-r+1$ from it, and all cross slots of the later rounds

are omitted in both executions. For the early rounds this distance is almost the length d of the main line. So the changed set and its witness must sit at the two ends of the line, while the chain has to change every set. The chain (2) changes half of the line at a time, with witnesses at distance only L, and so gives only $R \geq L + 1$. We therefore place the sets on a cycle and move the place where we open it, so that every set is in turn an end of the line in use.

The proof first handles the cases $d \leq 2$. It then defines the cycle, the lines obtained by opening it, and the final pattern D^*, in which a party at the far end of the main line learns nothing about the sender's input through round R. Finally, it states the three conditions that every link satisfies, shows that the opening can be moved by such links (Claim 4.1), and builds the chain from the all-honest execution to E_{D^*} in three phases. They cut the rounds R down to 2, then the sender's round-1 messages to the far sets, and, if $n - t$ does not divide n, they split the last set of the cycle.

Theorem 1.2. *Let $t < n$. For any R-round t-secure UA-BC protocol, R is at least $\min\{n - 1, \lceil \frac{n}{n-t} \rceil + \lfloor \frac{n}{n-t} \rfloor - 1\}$, which is at least $\frac{2n}{n-t} - 2$ if $t \leq n - 2$.*

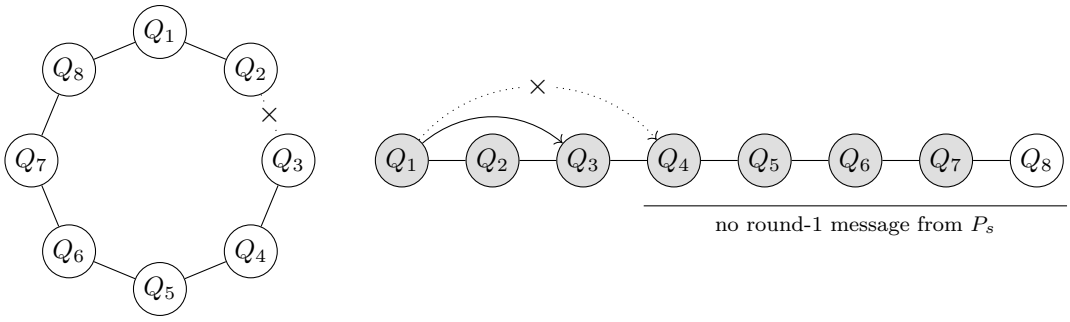
Proof. We use the probability space and notation of Section 3.1. Let R be the round complexity of the protocol. We show that $R \geq d$ for the line length $d = \min\{n-1, \lceil n/(n-t) \rceil + \lfloor n/(n-t) \rfloor - 1\}$ of Section 2, which is the bound of the theorem. For $t \leq n - 2$ the second term is the smaller one, and $\lceil x \rceil + \lfloor x \rfloor \geq 2x - 1$ gives $d \geq 2n/(n - t) - 2$.

The cases $d \leq 2$. In the all-honest execution, a party other than the sender outputs B by round R except with negligible probability, while its view before the first round is independent of B, so $R \geq 1$. This proves the case $d = 1$. We derive the case $d = 2$ from the chain (2) with $k = 1$, since the cycle below needs at least four sets. As $d = 2$, we have $n \geq 3$, $t \geq 1$ and $L = 1$, so the main line has the three sets $Q_1 = \{s\}$, Q_2 and Q_3. Suppose that $R = 1$ and fix $w \in Q_3$. By Lemma 3.2, except with negligible probability, termination within one round and validity hold in $E_\emptyset$, and termination within one round and agreement hold against $\mathcal{A}_{D'_1,1}$ and $\mathcal{A}_{D'_1,2}$. Then every party outputs B by round 1 in $E_\emptyset$, and hence in E_{D_2}, which delivers the same messages in round 1. By Lemma 3.1, in its first use in Section 3.1 with $\rho' = \infty$ and $J = \{2, 3\}$, the sender has the same view through round $1 + L - 1 = 1$ in E_{D_2} and in $E_{D'_1}$, so it outputs B by round 1 in $E_{D'_1}$, and by agreement so does P_w. The lemma with $J = \{1\}$ shows that P_w has the same view through round $1 + d - L = 2$ in $E_{D'_1}$ and in E_{D_1}, so it outputs B by round 1 in E_{D_1}. Hence $\Pr[Y_w(D_1) = B \wedge T_w(D_1) \leq 1] \geq 1 - \mathsf{negl}(\kappa)$, which contradicts the bound (1) with $r = 1 = d - 1$ for all sufficiently large κ. So let $d \geq 3$ and suppose that $1 \leq R \leq d - 1$.

The cycle of sets. We arrange sets $Q_1, \ldots, Q_m$ on a cycle, with indices modulo m, such that $Q_1 = \{s\}$ and $|Q_i \cup Q_{i+1}| \geq h$ for all i. For $h = 1$ they are the n singletons, $m = n = d + 1$ and $e \triangleq 0$. For $h \geq 2$ we write $n = Lh + e$ with $0 \leq e < h$, where $L \geq 2$ as $d \geq 3$, and let $m \triangleq 2L$, with sizes $1, h - 1, 1, h - 1, \ldots$ and the last set of size $h - 1 + e$. For $e = 0$ these are exactly the sets of the main line of Section 3.1. For $e > 0$ they are its sets $Q_1, \ldots, Q_{m-1}$ followed by the union of its last two sets, which we denote Q_m in this section and split again in the final line below. Then $d = m - 1$ if $e = 0$ and $d = m$ if $e > 0$, so that $m - 1 \geq R$ in both cases, as $R \leq d - 1$.

The lines obtained by opening the cycle. For $j \in \mathbb{Z}_m$, opening the cycle between Q_j and Q_{j+1} gives the line $\Lambda_j \triangleq (Q_{j+1}, Q_{j+2}, \ldots, Q_{j+m})$, whose end sets Q_{j+1} and Q_j are at distance $m-1$ (Figure 8(a)). Let X be the set of slots of rounds 1 to R between sets at cyclic distance at least 2, that is, between sets that are not adjacent on the cycle. As $m \geq 4$, the Λ_j-cross slots of these rounds are those in $X_j \triangleq X \cup \{\text{slots of rounds 1 to } R \text{ between } Q_j \text{ and } Q_{j+1}\}$. The slots in X are Λ_j-cross slots for every j, and X_j adds the slots across the current opening.

Finally, the blindness argument below needs a line whose far end is at distance d from the sender, that is, a line with $d + 1$ sets, while the ends of every Λ_j are at distance $m - 1$, which is $d - 1$ when $e > 0$. If $e = 0$ we have $m - 1 = d$ and take $\Lambda^* \triangleq \Lambda_m = (Q_1, \ldots, Q_m)$. If $e > 0$, the last



(a) (b)

Figure 8: (a) The cycle for $m = 8$, with $Q_1 = \{s\}$. When the slots of X_2 are omitted, messages are delivered only within a set and between sets joined by a solid edge, which form the line $\Lambda_2 = (Q_3, \ldots, Q_8, Q_1, Q_2)$. (b) The final omission pattern D^* on $\Lambda^* = \Lambda_8$ for $m = 8$, so that $d = 7$, and $R = 5$. The sender's round-1 messages reach Q_2 and Q_3 but no set from Q_4 on, and from round 2 on only adjacent sets communicate. So by round 5 information about B can reach the shaded sets but not Q_8. The message to Q_3 need not be omitted, since Q_3 is at distance $R = 5$ from Q_8, so what it learns in round 1 reaches Q_8 only in round 6.

set Q_m has $h - 1 + e$ parties, and we split it into Q' with $h - 1$ parties and Q'' with the remaining $e \geq 1$ parties, the last two sets of the main line, and let $\Lambda^* \triangleq (Q_1, \ldots, Q_{m-1}, Q', Q'')$. This is a line, as $|Q_{m-1} \cup Q'| = h$ and $|Q' \cup Q''| = h - 1 + e \geq h$. Splitting the set on the cycle instead may violate the requirement on adjacent sets, since there Q'' would also be adjacent to $Q_1 = \{s\}$ and $|Q'' \cup Q_1| = e + 1$ may be smaller than h. In both cases Λ^* is the main line of Section 3.1, with $d + 1$ sets, which we number by positions 1 to $d + 1$, starting with Q_1, so that for $e > 0$ the sets Q' and Q'' are at positions d and $d + 1$.

The final omission pattern D^*. Let D^* consist of all Λ^*-cross slots of rounds 2 to R and of the round-1 slots from P_s to the parties at positions $d + 2 - R$ to $d + 1$ of Λ^* (Figure 8(b)). In E_{D^*}, from round 2 on, messages travel only within a set or between adjacent sets of Λ^*, and the sender's round-1 messages reach only the positions up to $d + 1 - R$. For a fixed setup σ and tapes u, compare $E_{D^*}(0, \sigma, u)$ and $E_{D^*}(1, \sigma, u)$. Only the sender's messages depend on its input, so at the end of round 1 only parties at positions up to $d + 1 - R$ can have different views, and by Lemma 3.1 with $\rho = 1$ and $\rho' = R$ a party P_w in the set at position $d + 1$ has the same view through round R in both. So the view of P_w through round R in E_{D^*} is a function of (Σ, U) alone, independent of B, and as in (1) we get $\Pr[Y_w(D^*) = B \wedge T_w(D^*) \leq R] \leq \frac{1}{2}$. The theorem follows once we show that the guarantees of π force $Y_w(D^*) = B$ and $T_w(D^*) \leq R$ except with negligible probability, that is, force P_w to output B by round R in E_{D^*}, which we do with a chain of links from $\emptyset$ to D^* that each keep the common output. As in Section 3.1, the chain cuts the rounds backward, from R down to 2, because a link that changes round r needs all cross slots of the later rounds omitted already. Unlike there, it cuts each round set by set rather than in two halves. For round r it adds to the pattern all slots between sets that are not adjacent on the cycle, together with those across the current opening, rotating the opening so that each set is an end of the current line when the slots into it are added. Round 1 comes last and is treated differently, since only the sender's messages carry B. There the chain adds only the round-1 slots from the sender to the far sets of Λ^*, and if $e > 0$ it finally splits the last set.

Links and witnesses. Every omission pattern in the chain below contains only slots of rounds 1 to R, since with deterministic termination every party outputs by round R and a message of a

later round cannot affect an output made by then. Each link of the chain compares the executions E_D and $E_{D'}$ of two such patterns D and D', and satisfies the following three conditions for some line Λ and some round $r \in [R]$.

1. Every slot in D and in D', in every round, is a Λ-cross slot.
2. Both patterns contain all Λ-cross slots of rounds $r+1, \ldots, R$.
3. D and D' differ only in slots of round r, and the receivers of all these slots lie in a single set of Λ, the *changed set*.

Condition 1 makes $\mathcal{A}^{\Lambda}_{D,q}$ and $\mathcal{A}^{\Lambda}_{D',q}$ legal attacks for every q, so by Section 3.1 all parties have the same output $Y(D)$ in E_D, and all parties have the same output $Y(D')$ in $E_{D'}$, on the event that agreement and termination hold against these attacks. Showing that $Y(D) = Y(D')$ there is the purpose of the link. A *witness* of the link is a party in a set of Λ at distance at least $R - r + 1$ from the changed set. By condition 3, E_D and $E_{D'}$ deliver the same messages before round r and differ in round r only in messages into the changed set, and by condition 2 Lemma 3.1 applies with $\rho = r$, $\rho' = R$ and J the position of the changed set, so the witness has the same view through round R in E_D and in $E_{D'}$. On every value of Coins on which it outputs by round R, which deterministic termination guarantees except with negligible probability, it outputs the same value at the same round in both executions, and $Y(D) = Y(D')$. A link here thus plays the same role as a link of the chain (2), with two differences. It may change the slots into any set of the line, as long as the witness is far enough from that set, and its witness must keep its view through the round bound R rather than through L.

The next claim moves the opening of the cycle one set further in all rounds after ρ, by links alone and without touching the rounds up to ρ. It is the single move that the phases below repeat.

Claim 4.1. *Let $\rho \in [R]$ and $j \in \mathbb{Z}_m$, and let D be an omission pattern of rounds 1 to R whose slots of rounds up to ρ lie in X and whose slots of rounds $\rho+1, \ldots, R$ are exactly those of X_j. Then there is a sequence of omission patterns, starting with D and ending with the pattern that agrees with D on rounds up to ρ and whose slots of rounds $\rho+1, \ldots, R$ are exactly those of X_{j+1}, in which every two consecutive patterns form a link with the three conditions above.*

Proof. The sequence has two halves. In the first half we remove, for $r = \rho+1, \ldots, R$ in increasing order, the round-r slots between Q_j and Q_{j+1}, first those into Q_j and then those into Q_{j+1}. After this half, the rounds $\rho+1, \ldots, R$ omit exactly the slots of X, as if the cycle had no opening. In the second half we add, for $r = R, \ldots, \rho+1$ in decreasing order, the round-r slots between Q_{j+1} and Q_{j+2}, first those into Q_{j+1} and then those into Q_{j+2}, which leaves exactly the slots of X_{j+1} in these rounds. Every link of the first half is a link on Λ_j. All slots of both of its patterns lie in X_j, the slots of the rounds after r are still those of X_j, and the changed set is an end of Λ_j, so a party at the other end, in Q_{j+1} when Q_j changes and in Q_j when Q_{j+1} changes, is a witness at distance $m - 1 \geq R - r + 1$. Every link of the second half is a link on Λ_{j+1}, whose ends are Q_{j+2} and Q_{j+1}, for the same reasons: all slots lie in X_{j+1}, the slots of the rounds after r are already those of X_{j+1}, and the witness is a party at the end opposite to the changed set. The increasing order in the first half and the decreasing order in the second are what keep every round after the changed one omitting all cross slots of the line in use. □

Phase 1, cutting the rounds R to 2. We now build the chain from $\emptyset$, starting with $j = m$. The first phase makes every round from R down to 2, one round at a time in this order, omit exactly the cross slots of the line in use, so that in these rounds messages travel only within a set or between adjacent sets of that line, and it leaves round 1 untouched. When it reaches round ρ, the pattern has no slots in rounds up to ρ and exactly the slots of X_j in rounds $\rho+1, \ldots, R$, for

the current j. For round ρ it repeats the following operation m times. It adds the round-ρ slots of X into Q_j, then those into Q_{j+1}, and then moves the opening by one set with Claim 4.1, which increases j by one. After the m repetitions it adds the round-ρ slots across the current opening, between Q_j and Q_{j+1}, first those into Q_j and then those into Q_{j+1}.

We verify this. The two additions of each repetition, and the two at the end, are links on Λ_j: all slots of the pattern lie in X_j, the rounds after ρ contain all slots of X_j, and the changed set is an end of Λ_j, so a party at the other end, in Q_{j+1} when Q_j changes and in Q_j when Q_{j+1} changes, is a witness at distance $m-1 \geq R-\rho+1$. The claim applies at each repetition, since the slots of rounds up to ρ lie in X. After the m repetitions, j is back at its initial value and every set has been Q_j once, so round ρ contains all round-ρ slots of X, and the two final links make the round-ρ slots of the pattern exactly those of X_j, so the pattern is ready for round $\rho-1$. A round-ρ slot of X into a set can be added only while that set is an end of the line in use, which is why the opening travels around the whole cycle once for every round.

Phase 2, cutting the sender's round-1 messages. At the start of the second phase, the slots of rounds 2 to R are those of X_j for the current j, and round 1 has no slots. The phase adds the round-1 slots from P_s to the sets at positions $d+2-R$ to $d+1$ of Λ^*, and nothing else in round 1. Only the sender's round-1 messages depend on B, and these positions are the ones within distance $R-1$ of position $d+1$, from which B would reach position $d+1$ by round R. For $a = d+2-R, \ldots, m-1$ in increasing order, where $d+2-R \geq 3$ as $R \leq d-1$, we apply Claim 4.1 with $\rho = 1$ until $j = a$, so that Q_a is an end of the line in use Λ_a, and add the round-1 slots from P_s into Q_a. This is a link on Λ_a with a witness in Q_{a+1}, the other end of Λ_a, at distance $m-1 \geq R$ from Q_a. The slots added so far in round 1 lie in X, since $3 \leq a \leq m-1$ means that Q_1 and Q_a are not adjacent on the cycle, so the hypothesis of the claim on round 1 keeps holding. Finally we apply the claim until $j = m$ and add the round-1 slots from P_s to the parties of Q_m at positions $d+2-R$ to $d+1$ of Λ^*. These are all of Q_m, except that for $e > 0$ and $R = 1$ they are only Q''. This is a link on Λ_m, whose ends are Q_1 and Q_m, with the witness P_s at distance $m-1 \geq R$.

Phase 3, splitting the last set. If $e > 0$, the pattern now omits in rounds 2 to R the slots of X_m, the Λ_m-cross slots of these rounds, while D^* omits the Λ^*-cross slots. The line Λ^* arises from Λ_m by splitting Q_m into Q' and Q'', which makes Q_{m-1} and Q'' non-adjacent, so its cross slots are those of X_m together with the slots between Q_{m-1} and Q''. A third phase adds these slots for the rounds $r = R, \ldots, 2$ in decreasing order, first those into Q'' and then those into Q_{m-1}. Both are links on Λ^*: all slots of the pattern are Λ^*-cross slots, the rounds after r already contain all of them, and the witness P_s is at distance m from Q'' and at distance $m-2 \geq R-r+1$ from Q_{m-1}, as $r \geq 2$ and $R \leq d-1 = m-1$.

The chain ends at D^*. After the three phases, the pattern omits all Λ^*-cross slots of rounds 2 to R and the round-1 slots from P_s to the positions $d+2-R$ to $d+1$, and nothing else, that is, it is D^*. Each application of Claim 4.1 consists of at most $4R$ links, and the phases apply it at most Rm times and add at most $4Rm$ further links, so the chain has $O(d^3)$ patterns. Each of them consists of Λ-cross slots for a line Λ with at most $d+1$ sets, so the chain involves polynomially many adversaries $\mathcal{A}^{\Lambda}_{D,q}$, which can be computed from n, t and R. Let Good be the event that termination within R rounds and agreement hold against all of them and validity holds in $E_\emptyset$. By Lemma 3.2, $\Pr[\mathsf{Good}] \geq 1 - \mathsf{negl}(\kappa)$. Fix a value $\mathsf{coins} = (b, \sigma, u)$ of Coins in Good. Then all parties output b in $E_\emptyset$, every execution of the chain has a common output, and every link keeps it, since its witness outputs by round R with the same view in both executions. By induction along the chain, $y(D^*) = b$, and P_w outputs by round R. So $\mathsf{Good} \subseteq \{Y_w(D^*) = B \wedge T_w(D^*) \leq R\}$, and this event has probability at least $1 - \mathsf{negl}(\kappa)$, which exceeds the bound $\frac{1}{2}$ established for D^* above for all sufficiently large κ, a contradiction. □

The proof uses termination and agreement only against oblivious omission adversaries and validity only in the all-honest execution, so Theorem 1.2 also holds for broadcast with weak validity, where the sender's input is required only when all parties are honest. Deterministic termination only bounds the round by which every honest party terminates, so the theorem also holds when the honest parties terminate in different rounds. For broadcast, the single-execution argument of [GKKO07] on the main line, which Section 3 recalls, gives the same bound d, but it needs an honest sender at one end of the line.

For protocols with perfect guarantees, the chain argument of [DS83] gives $\min\{t+1, n-1\}$ rounds, and it applies to UA-broadcast since its endpoints are all-honest executions, see also [Du89]. Once a negligible error is allowed, the chains of Chor, Merritt and Shmoys [CMS89] and of Karlin and Yao [KY84] bound the error probability of an R-round protocol from below in terms of R, which rules out a constant R when $t = \omega(1)$. They apply to UA-broadcast [CCF+26], but they do not depend on $n/(n-t)$. So each classical lower bound for broadcast with deterministic termination, namely those of [DS83], of [CMS89, KY84] and of [GKKO07], also holds for UA-broadcast, and the known lower bounds for broadcast and for UA-broadcast differ only in expectation, where the line argument gives d for broadcast in a single execution, while Theorem 1.1 gives $t/(2(n-t))$ for UA-broadcast against rushing omission adversaries.

Tightness. The protocol of Dolev and Strong [DS83], which terminates deterministically in $t+1$ rounds, matches the bound of Theorem 1.2 exactly for $t \geq n-2$, since its variant for $n-2$ corruptions in which the sender outputs its own input takes $n-1 = d$ rounds and also tolerates $n-1$ corruptions. For every constant number $h = n-t$ of honest parties it takes $\Theta(n) = \Theta(d)$ rounds, so the bound is asymptotically tight for $h = O(1)$.

AI Methodology

We used Claude (Anthropic) to assist with the writing of this paper. We used it to revise and polish the text of all sections, to draw the figures and to format Table 1. We also used it to check the correctness of all our proofs. The authors verified the correctness and originality of all content, including references.

Bibliography

[Asp12] James Aspnes. Faster randomized consensus with an oblivious adversary. In Darek Kowalski and Alessandro Panconesi, editors, *31st ACM PODC*, pages 1–8. ACM, July 2012.

[BE03] Michael Ben-Or and Ran El-Yaniv. Resilient-optimal interactive consistency in constant time. *Distributed Computing*, 16(4):249–262, 2003.

[Ben83] Michael Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols (extended abstract). In Robert L. Probert, Nancy A. Lynch, and Nicola Santoro, editors, *2nd ACM PODC*, pages 27–30. ACM, August 1983.

[BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). In *20th ACM STOC*, pages 1–10. ACM Press, May 1988.

[CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). In *20th ACM STOC*, pages 11–19. ACM Press, May 1988.

[CCF+26] Ran Cohen, Daniel Collins, Pouyan Forghani, Juan Garay, and Vassilis Zikas. On the round complexity of dishonest-majority MPC. In *TCC 2026*, LNCS. Springer, Cham, 2026.

[CCGZ19] Ran Cohen, Sandro Coretti, Juan A. Garay, and Vassilis Zikas. Probabilistic termination and composability of cryptographic protocols. *Journal of Cryptology*, 32(3):690–741, July 2019.

[CCGZ21] Ran Cohen, Sandro Coretti, Juan A. Garay, and Vassilis Zikas. Round-preserving parallel composition of probabilistic-termination cryptographic protocols. *Journal of Cryptology*, 34(2):12:1–12:57, March 2021.

[CDL+25] Ran Cohen, Jack Doerner, Eysa Lee, Anna Lysyanskaya, and Lawrence Roy. An unstoppable ideal functionality for signatures and a modular analysis of the dolev-strong broadcast. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part IV*, volume 16271 of *LNCS*, pages 675–708. Springer, Cham, December 2025.

[CEG26] Pierre Civit, Grigorii Emdin, and Rachid Guerraoui. Constant-round interactive consistency in dishonest majority, or with quadratic communication in honest majority. Cryptology ePrint Archive, Paper 2026/1146, 2026.

[CFG+23] Ran Cohen, Pouyan Forghani, Juan A. Garay, Rutvik Patel, and Vassilis Zikas. Concurrent asynchronous byzantine agreement in expected-constant rounds, revisited. In Guy N. Rothblum and Hoeteck Wee, editors, *TCC 2023, Part IV*, volume 14372 of *LNCS*, pages 422–451. Springer, Cham, November / December 2023.

[CGZ20] Ran Cohen, Juan A. Garay, and Vassilis Zikas. Broadcast-optimal two-round MPC. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part II*, volume 12106 of *LNCS*, pages 828–858. Springer, Cham, May 2020.

[Cle86] Richard Cleve. Limits on the security of coin flips when half the processors are faulty (extended abstract). In *18th ACM STOC*, pages 364–369. ACM Press, May 1986.

[CMS89] Benny Chor, Michael Merritt, and David B. Shmoys. Simple constant-time consensus protocols in realistic failure models. *Journal of the ACM*, 36(3):591–614, 1989.

[CPS20] T.-H. Hubert Chan, Rafael Pass, and Elaine Shi. Sublinear-round byzantine agreement under corrupt majority. In Aggelos Kiayias, Markulf Kohlweiss, Petros Wallden, and Vassilis Zikas, editors, *PKC 2020, Part II*, volume 12111 of *LNCS*, pages 246–265. Springer, Cham, May 2020.

[DLM82] Richard A. DeMillo, Nancy A. Lynch, and Michael Merritt. Cryptographic protocols. In *14th ACM STOC*, pages 383–400. ACM Press, May 1982.

[DS83] Danny Dolev and H. Raymond Strong. Authenticated algorithms for Byzantine agreement. *SIAM Journal on Computing*, 12(4):656–666, 1983.

[Du89] Dingzhu Du. Lower bounds for weak byzantine agreement. *Acta Mathematicae Applicatae Sinica*, 5:355–360, 1989.

[FGH+02] Matthias Fitzi, Daniel Gottesman, Martin Hirt, Thomas Holenstein, and Adam Smith. Detectable byzantine agreement secure against faulty majorities. In Aleta Ricciardi, editor, *21st ACM PODC*, pages 118–126. ACM, July 2002.

[FGM01] Matthias Fitzi, Nicolas Gisin, and Ueli Maurer. Quantum solution to the Byzantine agreement problem. *Physical Review Letters*, 87(21):217901, 2001.

[FGMvR02] Matthias Fitzi, Nicolas Gisin, Ueli M. Maurer, and Oliver von Rotz. Unconditional byzantine agreement and multi-party computation secure against dishonest minorities from scratch. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 482–501. Springer, Berlin, Heidelberg, April / May 2002.

[FL82] Michael J. Fischer and Nancy A. Lynch. A lower bound for the time to assure interactive consistency. *Information Processing Letters*, 14(4):183–186, 1982.

[FM88] Paul Feldman and Silvio Micali. Optimal algorithms for byzantine agreement. In *20th ACM STOC*, pages 148–161. ACM Press, May 1988.

[FN09] Matthias Fitzi and Jesper Buus Nielsen. On the number of synchronous rounds required for byzantine agreement. In *International Symposium on Distributed Computing (DISC)*, pages 449–463, 2009.

[GGLZ22] Diana Ghinea, Vipul Goyal, and Chen-Da Liu-Zhang. Round-optimal byzantine agreement. In Orr Dunkelman and Stefan Dziembowski, editors, *EUROCRYPT 2022, Part I*, volume 13275 of *LNCS*, pages 96–119. Springer, Cham, May / June 2022.

[GKKO07] Juan A. Garay, Jonathan Katz, Chiu-Yuen Koo, and Rafail Ostrovsky. Round complexity of authenticated broadcast with a dishonest majority. In *48th FOCS*, pages 658–668. IEEE Computer Society Press, October 2007.

[GL05] Shafi Goldwasser and Yehuda Lindell. Secure multi-party computation without agreement. *Journal of Cryptology*, 18(3):247–287, July 2005.

[GMW87] Oded Goldreich, Silvio Micali, and Avi Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In Alfred Aho, editor, *19th ACM STOC*, pages 218–229. ACM Press, May 1987.

[Had84] Vassos Hadzilacos. *Issues of Fault Tolerance in Concurrent Computations*. PhD thesis, Harvard University, 1984.

[IOZ14] Yuval Ishai, Rafail Ostrovsky, and Vassilis Zikas. Secure multi-party computation with identifiable abort. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part II*, volume 8617 of *LNCS*, pages 369–386. Springer, Berlin, Heidelberg, August 2014.

[KK06] Jonathan Katz and Chiu-Yuen Koo. On expected constant-round protocols for byzantine agreement. In Cynthia Dwork, editor, *CRYPTO 2006*, volume 4117 of *LNCS*, pages 445–462. Springer, Berlin, Heidelberg, August 2006.

[KY84] Anna R. Karlin and Andrew Chi-Chih Yao. Probabilistic lower bounds for Byzantine agreement and clock synchronization. Unpublished manuscript, 1984.

[LLR02] Yehuda Lindell, Anna Lysyanskaya, and Tal Rabin. Sequential composition of protocols without simultaneous termination. In Aleta Ricciardi, editor, *21st ACM PODC*, pages 203–212. ACM, July 2002.

[LSP82] Leslie Lamport, Robert E. Shostak, and Marshall C. Pease. The Byzantine generals problem. *ACM Trans. Program. Lang. Syst.*, 4(3):382–401, 1982.

[PSL80] Marshall C. Pease, Robert E. Shostak, and Leslie Lamport. Reaching agreement in the presence of faults. *Journal of the ACM*, 27(2):228–234, 1980.

[PT86] Kenneth J. Perry and Sam Toueg. Distributed agreement in the presence of processor and communication faults. *IEEE Transactions on Software Engineering*, 12(3):477–482, 1986.

[Rab83] Michael O. Rabin. Randomized byzantine generals. In *24th FOCS*, pages 403–409. IEEE Computer Society Press, November 1983.

[RB89] Tal Rabin and Michael Ben-Or. Verifiable secret sharing and multiparty protocols with honest majority (extended abstract). In *21st ACM STOC*, pages 73–85. ACM Press, May 1989.

[WXSD25] Jun Wan, Hanshen Xiao, Elaine Shi, and Srinivas Devadas. Expected constant round Byzantine broadcast under dishonest majority. *Journal of the ACM*, 72(4):30:1–30:39, 2025.

[ZHM09] Vassilis Zikas, Sarah Hauser, and Ueli M. Maurer. Realistic failures in secure multi-party computation. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 274–293. Springer, Berlin, Heidelberg, March 2009.

A The Lower Bound against Oblivious Adversaries

Section 3.2 lets the adversary choose its omissions during the execution. This appendix shows what the chain (2) gives against the oblivious omission adversaries of Section 2.1, which fix their omissions before Coins is sampled. The expected round complexity against these adversaries is the supremum of Section 2.2 taken only over oblivious omission t-adversaries. Theorem 1.1, whose adversaries are rushing, gives no bound on it, so Theorem A.1 below is not subsumed by it.

Theorem A.1. *Let $t \le n-2$ and let π be any t-secure UA-BC protocol, then its expected round complexity against oblivious omission adversaries is at least $\frac{1}{6}\sqrt{\frac{n}{n-t}} - 1$.*

Proof. The theorem follows from a trade-off between the speed of the all-honest execution and the speed under attack by an oblivious omission t-adversary, which we state and prove as a lemma before completing the proof.

Lemma A.2. *Let $n \ge 3$, $t \ge 1$ and $L = \lfloor (d+1)/2 \rfloor$, and let π be any t-secure UA-BC protocol. There is a negligible function ν such that for every $c \in \{0, \dots, L\}$ and $\varepsilon \ge 0$, if in the all-honest execution some party outputs by round c with probability at least $1-\varepsilon$, then against some oblivious omission t-adversary some honest party has not output by round L with probability at least $(1-2\varepsilon)/(4c+1) - \nu(\kappa)$, and so has expected output round at least $(L+1)(1-2\varepsilon)/(4c+1) - \nu(\kappa)$.*

Proof of Lemma A.2. We use the probability space and notation of Section 3.1. The hypotheses $n \ge 3$ and $t \ge 1$ give $d \ge 2$, as $n-1 \ge 2$ and $n/(n-t) > 1$, so the main line $(Q_1, \dots, Q_{d+1})$ of Section 3.1 has at least three sets. We will apply the blindness bound (1) to a party in its last set Q_{d+1} at round L, which needs $L \le d-1$. This holds since $L = \lfloor (d+1)/2 \rfloor \le d-1$ for $d \ge 3$ and $L = 1 = d-1$ for $d = 2$. We fix $w \in Q_{d+1}$.

The chain of executions. We walk from the all-honest execution to the line execution through the chain

$$E_\emptyset \longrightarrow E_{D_{c+1}} \longrightarrow E_{D'_c} \longrightarrow E_{D_c} \longrightarrow \cdots \longrightarrow E_{D_2} \longrightarrow E_{D'_1} \longrightarrow E_{D_1},$$

where the patterns D'_ρ are those of Section 3.1. The first link omits the cross slots of all rounds after c at once. By hypothesis, with probability at least $1-\varepsilon$ some party outputs by round c in $E_\emptyset$, and $E_{D_{c+1}}$ delivers the same messages as $E_\emptyset$ in the first c rounds, so this output is preserved. The other $2c$ links omit the cross slots of the rounds $c, c-1, \dots, 1$, one round at a time, and each round in two halves, first into the right half and then into the left half, as in the chain (2). Figures 5 and 6 show these links for seven sets and $c = 2$.

The attacks and the good event. Let $\mathcal{D} \triangleq \{D_1, \dots, D_{L+1}, D'_1, \dots, D'_L\}$, which contains the patterns of the chain for every $c \le L$. We consider the $N \triangleq (2L+1)d + 1 \le n^2$ adversaries $\mathcal{A}_{D,q}$ with $D \in \mathcal{D}$ and $q \in [d]$, together with the adversary that corrupts no party. So each execution E_D with $D \in \mathcal{D}$ is at the same time the execution under the d attacks that leave one adjacent pair of sets honest, as in Section 3.1. Let Good be the event that agreement and termination hold against every $\mathcal{A}_{D,q}$ with $D \in \mathcal{D}$ and validity holds in $E_\emptyset$. By Lemma 3.2, Good fails only with a negligible

probability δ. As the family $\mathcal{D}$, and hence the list of adversaries, is the same for every c, this δ does not depend on c, which is what makes the function ν of the lemma independent of c. On Good, for every $D \in \mathcal{D}$, every party outputs in E_D and all parties have the same output $Y(D)$, and every party outputs B in $E_\emptyset$.

Preservation of the common output. Let p be the largest of the probabilities $\Pr[T_s(D_{\rho+1}) > L]$ and $\Pr\left[T_w(D'_\rho) > L\right]$ for $\rho \in [c]$, and $\Pr[T_w(D_1) > L]$. Each of them is the probability that a party honest in one of our attacks has not output by round L, since the sender is honest against $\mathcal{A}_{D_{\rho+1},1}$, P_w is honest against $\mathcal{A}_{D'_\rho,d}$ and against $\mathcal{A}_{D_1,d}$, and the execution against $\mathcal{A}_{D,q}$ is E_D. Fix a value $\mathsf{coins} = (b, \sigma, u)$ of Coins in Good on which some party P_i outputs by round c in $E_\emptyset$, and $T_s(D_{\rho+1}) \le L$ and $T_w(D'_\rho) \le L$ for every $\rho \in [c]$. Until the union bound below, every execution is evaluated at coins. By validity P_i outputs b in $E_\emptyset$. It has the same view through round c in $E_{D_{c+1}}$, so it outputs b there as well, and $y(D_{c+1}) = b$. Now let $\rho \in [c]$ with $y(D_{\rho+1}) = b$. The patterns $D_{\rho+1}$ and D'_ρ differ only in cross slots of round ρ with receivers in the right half, and both contain all cross slots of the later rounds. So Lemma 3.1, in its first use in Section 3.1 with $\rho' = \infty$ and $J = \{L+1, \ldots, d+1\}$, shows that the sender, the *witness* of this link, has the same view through round $\rho + L - 1 \ge L$ in $E_{D_{\rho+1}}$ and in $E_{D'_\rho}$. As $T_s(D_{\rho+1}) \le L$, it outputs the same value at the same round in both, and $y(D'_\rho) = b$. Likewise, D'_ρ and D_ρ differ only in cross slots of round ρ with receivers in the left half, and the lemma with $J = [L]$ shows that P_w, the witness of this link, has the same view through round $\rho + d - L \ge L$ in $E_{D'_\rho}$ and in E_{D_ρ}, so $T_w(D'_\rho) \le L$ gives $y(D_\rho) = b$. Outputting by round L is sufficient for a witness, not necessary, since the horizons are $\rho + L - 1$ and $\rho + d - L$, but a uniform threshold keeps the accounting simple. By induction on $\rho = c, \ldots, 1$ we get $y(D_1) = b$, and in particular $y_w(D_1) = b$. So the event $Y_w(D_1) = B$ contains the intersection of Good, the event of the hypothesis and the $2c$ events $T_s(D_{\rho+1}) \le L$ and $T_w(D'_\rho) \le L$ for $\rho \in [c]$, and a union bound over their complements gives

$$\Pr[Y_w(D_1) = B] \;\ge\; 1 - \varepsilon - \delta - 2cp.$$

Blindness of the far endpoint. In the line execution P_w is in Q_{d+1} and $L \le d - 1$, so the blindness bound (1) with $r = L$ gives $\Pr[Y_w(D_1) = B \wedge T_w(D_1) \le L] \le \frac{1}{2}\Pr[T_w(D_1) \le L]$, and hence

$$\Pr[Y_w(D_1) = B] \;\le\; \tfrac{1}{2} + \tfrac{1}{2}\Pr[T_w(D_1) > L] \;\le\; \tfrac{1}{2} + \tfrac{p}{2}.$$

The chain gives $Y_w(D_1) = B$ whether or not P_w outputs by round L, but the blindness applies only when it does, which is why $\Pr[T_w(D_1) > L]$ enters with the factor $\frac{1}{2}$. The two bounds give $p(2c + \frac{1}{2}) \ge \frac{1}{2} - \varepsilon - \delta$, that is, $p \ge (1 - 2\varepsilon)/(4c + 1) - 2\delta/(4c + 1)$. As p is one of the listed probabilities, some party honest against one of the N adversaries has not output by round L with probability at least $(1 - 2\varepsilon)/(4c + 1) - 2\delta$, and its expected output round is at least $L + 1$ times this probability, that is, at least $(L + 1)(1 - 2\varepsilon)/(4c + 1) - 2(L + 1)\delta$. This proves the lemma with $\nu \triangleq 2(L + 1)\delta$, which is negligible as $L < n$, and which depends neither on c nor on ε. □

In the chain of Lemma A.2, which extends the chain (2) for $k = c$ by an all-honest execution in which some party outputs by round c, the probabilities that the witnesses of the $2c$ links are slow add up to at least about $\frac{1}{2}$. Each of them is the probability of a slow party in one fixed execution of the chain, which one oblivious omission adversary realizes, and the lemma can only take the largest of them, which may be as small as $1/(4c)$, since a protocol may have a slow party in every execution of the chain with the same probability. The adversaries of Section 3.2 choose their cut during the execution, and so realize on each value of Coins an execution of the chain in which some party is slow, if there is one, which makes their probability of a slow party the whole sum rather than its largest term.

The square root in Theorem A.1 comes from balancing the two sides of this trade-off. Let C_0 be the expected output round of the sender when all parties are honest. The adversary that corrupts no party gives the bound C_0. Lemma A.2 needs a round c by which the all-honest execution outputs with probability bounded away from $\frac{1}{2}$, and Markov's inequality shows that $c \approx 4C_0$ is such a round, with probability $\frac{3}{4}$, so the lemma gives a bound of about $L/(32C_0)$ under attack. Every protocol thus yields the larger of C_0 and about $L/(32C_0)$. This expression depends on the protocol only through C_0, and its minimum over all values of C_0, attained where the two sides meet, at C_0 of order $\sqrt{L}$, is the bound that holds uniformly for every protocol. The minimization says nothing about whether some protocol attains it. The rest of the proof carries this out, with $\sqrt{L}/6$ as a clean value slightly below the meeting point.

In the notation of Section 3.1, let B be the sender's input bit and $E_\emptyset$ the all-honest execution. If $d = 1$, the bound holds since a party other than the sender outputs B in $E_\emptyset$ except with negligible probability, while its view before the first round is independent of B, so its expected output round is at least $1 - \mathsf{negl}(\kappa)$. Let $d \geq 2$, so that $n \geq 3$ and $t \geq 1$, and let $C_0 \triangleq \mathbb{E}[T_s(\emptyset)]$ be the expected output round of the sender in $E_\emptyset$. If $C_0 \geq \sqrt{L}/6$, the adversary that corrupts no party and the sender give the bound. Otherwise $c \triangleq \max\{0, \lceil 4C_0 \rceil - 1\}$ satisfies $c \leq 4C_0 < \frac{2}{3}\sqrt{L}$, so $c \leq L-1$, and by Markov's inequality the sender outputs by round c with probability at least $3/4$. Lemma A.2 with $\varepsilon = \frac{1}{4}$ then gives an adversary and a party honest against it whose expected output round is at least $(L+1)/(8c+2) - \nu(\kappa)$, where ν is the negligible function of the lemma and does not depend on c. As $8c + 2 < \frac{16}{3}\sqrt{L} + 2$ and $6(L+1) > \frac{16}{3}L + 2\sqrt{L}$, this is at least $\sqrt{L}/6 - \nu(\kappa)$. In either case some party has expected output round at least $\sqrt{L}/6 - \mathsf{negl}(\kappa)$ while honest against an oblivious omission t-adversary, the sender against the adversary that corrupts no party in the first case, and a party under attack in the second. As $t \leq n-2$, we have $L = \lfloor n/(n-t) \rfloor \geq n/(n-t) - 1$, hence $\sqrt{L} \geq \sqrt{n/(n-t)} - 1$ and $\sqrt{L}/6 - \mathsf{negl}(\kappa) \geq \frac{1}{6}\sqrt{n/(n-t)} - \frac{1}{6} - \mathsf{negl}(\kappa)$, which is at least the bound of the theorem once $\mathsf{negl}(\kappa) \leq 5/6$. □

The proof uses $t \leq n-2$ only in its last sentence, and for $t = n-1$ it gives $\sqrt{\lfloor n/2 \rfloor}/6 - \mathsf{negl}(\kappa)$, as then $L = \lfloor n/2 \rfloor$.

The proofs of this appendix use agreement and termination only against oblivious omission adversaries and validity only when all parties are honest, so their bounds hold for every primitive with these guarantees.

The remark in the proof of Theorem 1.1 on protocols that are polynomial-time in each round rather than in total applies here as well. The bound of Theorem A.1 survives in this case up to an additive $O(\sqrt{L}\,\kappa^{-\alpha})$ for any constant α. Let the corrupted parties of each $\mathcal{A}_{D,q}$ stop after round $P \triangleq n^3\kappa^\alpha$, which makes the adversaries polynomial-time and changes no view through round P, and let C be the largest expected output round of a party honest against one of the stopped adversaries. Then Good also requires every party to output by round P in every E_D with $D \in \mathcal{D}$, which by Markov's inequality adds at most $n(2L+1)C/(P+1)$ to δ. The party found in the proof of Theorem A.1 is honest against one of these adversaries, and whether it has output by round L is decided by round P, so the proof gives $C \geq \sqrt{L}/6 - 2(L+1)\delta - 2(L+1)n(2L+1)C/(P+1)$, that is, $C(1 + 4n^3/P) \geq \sqrt{L}/6 - 2(L+1)\delta$, as $(L+1)(2L+1) \leq 2n^2$.