More is Merrier: Relax the Non-Collusion Assumption in Multi-Server PIR

Tiantian Gong*, Ryan Henry†, Alexandros Psomas*, Aniket Kate*‡

*Purdue University *{tg, apsomas, aniket}@purdue.edu*
†University of Calgary *ryan.henry@ucalgary.ca*, ‡Supra Research

Abstract—A long line of research on secure computation has confirmed that anything that can be computed, can be computed securely using a set of non-colluding parties. Indeed, this non-collusion assumption makes a number of problems solvable, as well as reduces overheads and bypasses computational hardness results, and it is pervasive across different privacy-enhancing technologies. However, it remains highly susceptible to covert, undetectable collusion among computing parties. This work stems from an observation that if the number of available computing parties is much higher than the number of parties required to perform a secure computation task, collusion attempts in privacy-preserving computations could be deterred.

We focus on the prominent privacy-preserving computation task of multi-server 1-private information retrieval (PIR) that inherently assumes no pair-wise collusion. For PIR application scenarios, such as those for blockchain light clients, where the available servers can be plentiful, a single server's deviating action is not tremendously beneficial to itself. We can make deviations undesired via small amounts of rewards and penalties, thus significantly *raising the bar* for collusion resistance. We design and implement a collusion mitigation mechanism on a public bulletin board with payment execution functions, considering only rational and malicious parties with no honest non-colluding servers. Privacy protection is offered for an extended period after the query executions.

I. INTRODUCTION

The non-collusion assumption among computing parties is prevalent across privacy-preserving computations. It allows us to solve some important problems [23], [28], [61], with no known solutions without it. It also reduces overheads and allows us to bypass computational hardness results in other problems such as private information retrieval (PIR) [19]. In fact, it is fair to say that, the entire field of secure multi-party computation (MPC) [9], [93] relies on the non-collusion assumption. While most of these distributed systems do allow collusion among pre-defined subsets of parties, all bets are off once the adversary can lure a superset of parties. Indeed, using the non-collusion assumption remains highly susceptible to undetectable collusion among computing parties: collusion over Internet-based secure systems can be difficult to prove, if possible; parties may collude without recognizing each other. As these distributed systems get considered for applications such as privacy-preserving machine learning [15], [17], [56], [57], [69], [87], privacy-preserving database/directory searches [68], [75], and blockchain privacy [50], [77], the non-collusion assumption is considered an Achilles' heel for large adoption.

In the context of database querying service, a PIR protocol allows a client to query a database without revealing which piece of information is of interest. Chor et al. [19] formalize the k-server PIR problem ($k \geq 2$) and propose a mechanism for fetching a single bit from an N-bit string, which is stored by $\ell \geq k$ servers (or parties/players), without disclosing the bit's location to the k queried servers.

PIR can be constructed from a single server or multiple servers. Single server information-theoretic PIR (IT-PIR) is impossible beyond the trivial solution of sending the whole database to the user. The *inefficiency* of single-server computational PIR (cPIR) [60] schemes is also well-known. Consider a database $\mathbb{D}$ comprising N entries. The state-of-the-art proposals [2], [5], [25], [51], [67] usually have $O(\sqrt{N})$ communication complexity. OnionPIR [70] realizes $O(\log N)$ communication cost while incurring heavy computation costs. (Detailed efficiency comparisons reside in Appendix B.) By involving more *non-colluding* servers ($k \geq 2$), the communication and computation complexity can be substantially reduced, which serves as a motivation for multi-server PIR schemes [11], [18], [47], [81]. For example, in Hafiz-Henry's multi-server PIR [47], the query size is $130 \log k(\lceil \log N \rceil - 7)$ bits and the total response size has only $\frac{k}{k-1}$ times overhead. It also has the lowest possible computation cost [7]: for each bit being fetched, each server performs $\frac{|\mathbb{D}|}{k-1}$ bit operations on the database in expectation.

However, the non-collusion assumption promoting efficiency improvements calls for practical justification. Because servers can easily collude over unobserved side channels, or various covert and anonymous channels [83] to recover users' queried entries, and such **privacy-targeted collusion** remains oblivious to PIR procedures. The prevalence of these communication modalities makes the collusion problem impossible to solve within the PIR protocol.

In this paper, we find that by relaxing the non-collusion assumption to the rationality assumption on servers, deterring such privacy-targeted collusion becomes possible. Here, a rational party takes actions that maximize its utility. The core question is then *how to render collusion undesired for rational participants*, for which we draw intuition from a novel bit guessing game.

Secret-shared bit guessing game. An honest game host commits to a secret unbiased bit b and secret-shares it among k players. Here, every player receives a share such that bit b can be reconstructed with all k shares together. After sharing, each player can make *open* guesses about b and the game terminates when no more player wishes to guess. The *first*

correct guesser is selected as the winner. The host rewards the winner \$2 but takes \$4 from the other $k - 1$ players. Additionally, a player loses \$4 for a wrong guess. Without reconstructing bit b explicitly from received shares, rational players do not guess because the expected return is negative (\$-1). For players without a network advantage in submitting guesses first, not collaborating to reconstruct b is strictly dominating (i.e., generating strictly more returns) since their expected return is $(\frac{k-1}{k}(-4) + \frac{1}{k} \cdot 2)$ which is negative for $k \geq 2$. This thought experiment implies that it is possible to make unobserved collusion visible and undesired for rational participants via a proper incentive structure.

Deter collusion in multi-server PIR. The simple game above is already expressive, but there are caveats that can cause the incentives to stop being effective in practice, e.g., a rational host or bit b being biased. When designing a collusion-mitigation mechanism for multi-server PIR, we then face two challenges: (i) the difficulty of defining and verifying evidence of collusion when users and servers are not honest, and (ii) the intricate conflicts between ensuring the effectiveness of collusion mitigation and deterring manipulations from users and servers for biased and valuable secrets. Instead of making assumptions on collusion methods, we base the solution only on the premise that *at least some colluding parties learn user secrets or some nontrivial (i.e., output depends on the inputs) function of the secrets after successful collusion*. We then define collusion evidence to be outputs of non-trivial functions on user secrets and record commitments (e.g., cryptographic hash) of queries and responses for evidence verification.

For the incentives, users pay service fee λ_s for servers' PIR responses. We reward the colluding party (with λ_r) who reports collusion first and penalize the remaining accomplices by exerting penalty λ_p and depriving them of service fees. To avoid manipulations of the mechanism, we add randomness and apply fines λ_f to discourage servers with private knowledge about user secrets from making false accusations. We tune λ_s (when necessary) to hinder users' fake accusations. For parameterization of the payment amounts, we note the following:

In a single or known finite runs of the PIR service, we do not need many servers, i.e., $\ell = k$, or reporting rewards, i.e., $\lambda_r = 0$. But in the worst case, we need to calibrate the service fee and penalty amounts such that their weighted sum, $\lambda_s + \frac{k-1}{k}\lambda_p$, is greater than the secret's worth. This means that one expects to earn more by not colluding and earning the service fee. A larger k means looser condition on λ_s and λ_p, expanding parameter feasibility regions. Since $\lambda_r = 0$, servers and users do not falsely accuse when there are positive fines $\lambda_f > 0$ and service fees $\lambda_s > 0$.

In infinitely repeated runs of the service, we need to have many servers, i.e., $\ell \gg k$, and sufficiently positive reporting rewards, i.e., $\lambda_r > X$ where X is a positive term that decreases with ℓ (Theorem III.2). Intuitively, to achieve the non-collusion outcome in repeated runs, the same servers need to be queried less frequently to reduce the privacy worth they bear. And servers need to earn positive premiums by deviating from

an undesired strategy emerging in repeated games, which is always cooperating and not reporting. Since $\lambda_r > 0$, the fines λ_f and service fees λ_s now both need to exceed the reward by a certain amount to prevent false accusations.

A. Contributions

We devise a mechanism for relaxing the non-collusion assumption in multi-server PIR systems to rationality assumptions on servers. It is composed of *a practical incentive structure* and *a module for gathering and verifying collusion evidence*. Overall, we make the following contributions.

First and most importantly, we present a **practical collusion deterrence mechanism for multi-server PIR** so that one can enjoy the perfect efficiency of 1-private PIR.[1] (The mechanism directly applies to general t-private PIR that only assumes non-collusion among $> t$ parties.) This is achieved in harsh settings: (1) We let there be a secure strong collusion protocol at colluding parties' disposal. They can compute arbitrary functions on received queries or computed responses, e.g., reconstructing the queried entry. (2) Servers can hold *arbitrary* private knowledge about clients' queries. (3) Servers can collude over *any unobserved channels*. (4) A bounded number of servers are malicious. We employ many servers so that a single server does not bear too much information worth leaking before it is discouraged to collude (similar to when the share holders hesitate to recover bit b in the guessing game).

Second, we examine how servers exit the PIR service and adopt a privacy protection self-insurance to **protect user secrets for an extended period** after the queries. Third, we study how servers form strong coalitions, whose members have complete trust for each other and are immune to the provided incentives, from a cooperative game theory perspective. Our focus is identifying the conditions necessary to keep the coalition size small. Last, we **implement and deploy the mechanism** as a smart contract on Ethereum [34] blockchain and apply optimizations to economize the gas costs (in Table I).

Efficiency and practicality. Overall, the solution only adds to the underlying PIR scheme one signed commitment message for each query and response. The collusion resolution mechanism also has little overhead: it induces the non-collusion outcome among the rational servers and is not executed in equilibrium when there is no malicious server. For practicality, we observe that in many settings where the number of available computing parties can be much higher than the required set, an effective collusion mitigation mechanism with practical parameters, i.e., small fee λ_s and penalty λ_p, can be constructed.

Relevance to Blockchain Privacy. As a representative example, the proposed mechanism is useful for blockchain access privacy problem [50]: allowing light/thin clients that do not maintain the ledger locally to query full nodes in a privacy-preserving way [77]. Existing solutions are either inefficient

---

[1]1-private PIR schemes are constructed assuming no pairwise collusion and provide effectively perfect efficiency in terms of downloading capacity, in the sense that to fetch a b-bit entry, one only needs to download $b + 1$ bits [81] (by having exponentially many servers). Hafiz-Henry's construction [47] of 1-private multi-server PIR achieves the best possible download cost, $\frac{k}{k-1}Nb$ bits, with a fixed number of servers.

given the ever-changing nature of database [77] or require trusted hardware [62], [66], [91]. Blockchains are particularly suitable application scenarios for our mechanism as prominent blockchains have thousands of full nodes [10], [36], which admits more practical parameters for the design.

II. Preliminaries and overview

A. Definitions

Multi-server PIR. We follow the security definition for multi-server PIR by Hafiz-Henry [47]. Let I, Q, R, Y respectively denote the random variables for the entry index $i \in [N]$ of interest, the query string $q \in \{0, 1\}^*$, the response string $r \in \{0, 1\}^*$ and the reconstructed entry $y \in \mathbb{D}$.

Definition 1 (Multi-server PIR [47]). *The interaction described by (I, Q, R, Y) provides correctness if $\Pr[Y = \mathbb{D}_i | I = i] = 1$ where $\mathbb{D}_i$ is the i-th entry in $\mathbb{D}$. It provides perfect 1-privacy if $\forall i, i^* \in [N], \Pr[Q = q | I = i] = \Pr[Q = q | I = i^*]$ where the probability is over all the random coin tosses made by the client. It provides computational 1-privacy if $\forall i, i^* \in [N]$, the distribution ensembles $\{Q | I = i\}_{\lambda \in \mathbb{N}}$ and $\{Q | I = i^*\}_{\lambda \in \mathbb{N}}$ are computationally indistinguishable with λ being the security parameter.*

We focus on k-out-of-k 1-private PIR schemes where all k responses are needed in reconstruction. Note that the analysis applies to k'-out-of-k ($k' < k$) robust PIR schemes, but non-colluding parties can be penalized in the worst case (but not in equilibrium). We then denote the multi-server PIR scheme as $(\ell, k, 1)$-PIR. A set of ℓ servers $S = \{S_i | i \in [\ell]\}$ maintain identical databases consisting of N entries, $\mathbb{D} = \{\mathbb{D}_i | i \in [N]\}$. To query entry $\mathbb{D}_a$ at index a, user U generates k queries and sends them to k (distinct) servers. Each queried server locally computes responses to the query and responds with an answer string. U then collects all k response strings and reconstructs $\mathbb{D}_a$.

There is an abundance of available PIR constructions, including the ones for two servers [11], [19], [30], [78], [90], and three or more servers [4], [6], [18], [31], [47], [54], [59], [84], [94]. While all constructions suffice for this work, we find Boyle-Gilboa-Ishai's construction for 2-server cPIR [11] and Hafiz and Henry's construction for more server case [47] to be suitable. Both constructions realize PIR as a distributed point function and as elaborated in Appendix B, are highly efficient for million-sized databases.

Relevant concepts in game theory [71]. A normal-form game can be characterized with a set of players $\mathcal{P}$, actions available to players $\mathcal{A} = \times_{i \in \mathcal{P}} \mathcal{A}_i$, and utility functions $\mathcal{V} = \{u_1, \cdots, u_{|\mathcal{P}|}\}$ where $u_i : \mathcal{A} \mapsto \mathbb{R}$, mapping action profiles to real-numbered utilities. An extensive-form game further captures sequences of actions, beliefs over others' actions, knowledge of history, etc.

Let $\mathbf{s} \in \mathcal{A}$ be a strategy profile, which is simply a sample in universe $\mathcal{A}$. A *dominant strategy* $\mathbf{s}_i$ for a player $\mathcal{P}_i$ generates utility at least as good as any other strategies: $\forall \mathbf{s}_i' \neq \mathbf{s}_i, \forall \mathbf{s}_{-i}, u_i(\mathbf{s}_i, \mathbf{s}_{-i}) \geq u_i(\mathbf{s}_i', \mathbf{s}_{-i})$ where $\mathbf{s}_{-i}$ is the strategy profile for all other players. If equality never holds,

then $\mathbf{s}_i$ is a strictly dominant strategy. A *solution concept* then captures the equilibrium strategies for all players in a game. Our incentive structure induces a (extensive-form) sequential game and the solution concept we employ is subgame perfect equilibrium and sequential equilibrium, which we discuss more in Section III and Section A. Informally, in both equilibria, at each decision point of the sequential game, a player selects the action that maximizes the expected returns from that point on.

B. Threat model

We assume all servers are either rational, i.e., taking actions that maximize utilities, or malicious, i.e., behaving arbitrarily. We assume static corruption when faced with a malicious adversary. There are *no* honest non-colluding servers. Normally, servers provide PIR services by running a secure PIR protocol and are compensated by the system in the form of service fees. They can potentially collude with each other to break user privacy and obtain what the recovered private information is worth. Servers can collude as individuals (Section III), form mutually-trusted coalitions (Section V-B), or leave the service (Section V-A).

To make the results more general, we *handicap* ourselves by giving the adversary advantages in collusion. If these advantageous settings are not achieved in practice, the result still holds. First, servers may hold arbitrary knowledge about a user's index of interest. Second, servers can communicate with each other over any unobserved communication channels, e.g., anonymous channels that keep the communication covert and undetectable. Further, we assume the communication is simulatable so that a server can generate the communication transcripts on its own, and colluding servers can self-protect by denying having the conversation. Third, we let there exist a secure collusion protocol, e.g., a secure MPC protocol, that can be utilized by colluding parties to compute the result collectively.

C. Solution overview

We describe the solution given rational servers and present pseudocodes in Figure 1. Servers are deterministically indexed and participants all know the indices. We make use of a public bulletin board (**BB**) with payment execution functionality to instantiate the mechanism as a coordinator contract (**CC**). Algorithm **CC** verifies collusion evidence and executes payments involved in both normal service and collusion resolution. One implementation of **CC** is a smart contract on Ethereum [34] (Section VI-A).

Setup and normal service. Users know servers' identities, and they can directly communicate. To mitigate collusion, we have participants post commitments of messages on **BB** for potential evidence verification. Since we allow servers to hold arbitrary private knowledge, $\omega \geq 1$ random companion queries are sent to help distinguish between servers learning user secrets via collusion and via private knowledge.

We now describe the normal routines with one random companion query. Initially, as described in line 1 in Figure 1, a user samples one random index a_r and generates queries for

---

Normal Service (for querying $\mathbb{D}_a$)

1) User U samples an index a_r uniformly at random and runs two instances of $(\ell, k, 1)$-PIR for querying $\mathbb{D}_a$ and $\mathbb{D}_{a_r}$, generating $2k$ queries, $Q_1, \ldots, Q_k$ for $\mathbb{D}_a$ and $Q_1', \ldots, Q_k'$ for $\mathbb{D}_{a_r}$.

2) U commits to the $2k$ queries, obtaining $\mathbf{c} = (\langle c_1, \sigma_1 \rangle, \ldots, \langle c_k, \sigma_k \rangle)$ and $\mathbf{c}' = (\langle c_1', \sigma_1' \rangle, \ldots, \langle c_k', \sigma_k' \rangle)$. U picks a random permutation Π, posts $\Pi(\mathbf{c}, \mathbf{c}')$ to **BB**, and stores Π locally. Here, for $i \in [k]$, $c_i = \mathsf{Comm}(Q_i)$, $c_i' = \mathsf{Comm}(Q_i')$, $\sigma_i = \mathsf{Sign}(c_i)$ and $\sigma_i' = \mathsf{Sign}(c_i')$.

3) U samples k distinct integers in $[\ell]$ at random and obtains k corresponding server indices, $\mathbf{id} = (id_1, \ldots, id_k)$. U signs and posts $\langle \mathbf{id}, \sigma_{id} \rangle$ onto **BB**. Here, $\sigma_{id} = \mathsf{Sign}(\mathbf{id})$.

4) U sends corresponding de-commit information to each server S_{id_i} over a secure channel.

5) Each queried server S_{id_i} verifies and de-commits the i-th commitments, and computes the answers (A_i^1, A_i^2) to de-committed queries locally against database $\mathbb{D}$.

6) Each queried server S_{id_i} commits to the responses, signs the message and posts $(\langle c_{id_i}^1, \sigma_{id_i}^1 \rangle, \langle c_{id_i}^2, \sigma_{id_i}^2 \rangle)$ onto **BB**. S_{id_i} sends de-commit information to U over a secure channel. Here, $c_{id_i}^1 = \mathsf{Comm}(A_i^1)$, $c_{id_i}^2 = \mathsf{Comm}(A_i^2)$, $\sigma_{id_i}^1 = \mathsf{Sign}(c_{id_i}^1)$ and $\sigma_{id_i}^2 = \mathsf{Sign}(c_{id_i}^2)$.

7) U verifies and de-commits the commitments, retrieves responses, and reconstructs entry $\mathbb{D}_a$.

Collusion Resolution (Server S_{id_i} accusing)

8) S_{id_i} submits a collusion report to **CC**, indicating the accused parties and evidence, i.e., circuits for computing $f(\cdot)$, S_{id_i}'s inputs to $f(\cdot)$, and function outputs. **CC** rejects the report if $f(\cdot)$ is trivial.

9) Each accused server S_{id_j} provides auxiliary de-commit information for all possible input(s) to $f(\cdot)$.

10) **CC** confirms the accusation if (1) the accused server(s) fails to submit auxiliary information in time Δ, or (2) the de-commit info is incorrect, or the output of $f(\cdot)$ on de-committed inputs match ed.

Fig. 1: Overview of routines ($\omega = 1$ companion query) with payments omitted. $\mathsf{Comm}(\cdot)$ is a perfectly hiding and computationally binding commitment scheme and $\mathsf{Sign}(\cdot)$ is a secure digital signature scheme.

both the intended entry $\mathbb{D}_a$ and the random entry $\mathbb{D}_{a_r}$. The user then commits to the k pairs of queries, randomly permutes each pair, and posts the commitments to **BB** (line 2). Next, the user samples k servers randomly and send corresponding de-commit information to them (line 3-4). Queried servers then compute the responses (line 5), post commitments of responses to **BB** and send de-commit information to the user over any secure channels (line 6). The user then reconstructs $\mathbb{D}_a$ from the responses (line 7).

Collusion. Colluding servers can collectively execute any collusion protocol to learn some nontrivial function $f(\cdot)$ of the explicit bits of the index or entry (e.g., least significant bit) or learn other servers' de-committed queries or responses.

Collusion resolution. A collusion reporter provides the following as evidence: the circuit for computing function $f(\cdot)$, her inputs to $f(\cdot)$, and the output from computing $f(\cdot)$ (line 8 in Figure 1). Note that multiple inputs can be provided if functions on multiple secrets are output in an order unknown to the accuser. This does not change the analysis since we can deterministically verify the $(\omega + 1)$ possible outputs. We therefore focus on describing the generic resolution routine for a single input.

Upon receiving a report, the contract confirms the non-triviality of $f(\cdot)$ (details in Section VI-A), then counts down an auxiliary information collection window Δ. The accused parties who fail to provide such information to refute the evidence in time are treated as culpable. We discuss setting Δ

in Section VI-B. The accused server(s) open the corresponding commitments and reveal inputs to $f(\cdot)$ (line 9). **CC** verifies the correct opening of commitments and compares the output of $f(\cdot)$ with the evidence (line 10). Note that we do not distinguish between collusion that recovers the desired entry $\mathbb{D}_a$ and random entry $\mathbb{D}_{a_r}$.

Intuitions behind incentive design. The first correct collusion reporting server for a PIR run is selected as the winner. The payment rules are as follows: (1) fine servers submitting false reports λ_f; (2) distribute rewards λ_r to the winner; (3) penalize each server being successfully accused amount λ_p; (4) charge service fee λ_s for each query and pay a queried server that is not successful accused amount λ_s. (1) The framing fines are to discourage servers from false accusations. (2) Rewarding the first reporter is to encourage the servers participating in collusion to reveal its existence in infinitely repeated runs of the PIR service. (3) Fining the accused colluding parties is to discourage at least one party in the colluding group without a network advantage from colluding. (4) Collecting service fees from clients is first to stop clients from exploiting the mechanism and causing undue punishment on servers, and second to determine the length of the privacy protection period, which we discuss in more detail in Section V-A.

To realize the incentive scheme, in the beginning, servers make deposits greater than λ_p and λ_f to **CC**. Deposits function as sources of potential fines and penalties, and will be transferred back when servers exit the service. Users also

make deposits λ_s before making each query.

Organization of contents. We next dive into the key analysis for designing the mechanisms that achieve the non-collusion outcome while accounting for manipulations from servers and users in Section III. We add malicious servers in Section IV and update the collusion evidence verification routines. Section V discusses servers' exiting strategies, and how strong coalitions form. Section VI summarizes the mechanism and presents an implementation.

III. COLLUSION DETERRENCE

As mentioned in Section I, we base the analysis on the premise that after successful collusion, at least some colluders learn some nontrivial function of the user secrets. We now dissect the intrinsic incentive structure of the solution. We let all servers be rational and include malicious servers in Section IV.

A. Single-shot game

During a single-shot collusion, colluding servers verify the information input by others, e.g., by including circuits that examine users' signatures in de-commit information.

Order of events. Consider the following single-shot sequential game G:

(1) In round 1, servers decide whether to collude, C or $\bar{C}$.
(2) In round 2, colluding servers decide whether to input the correct information, amicable (A) or deceitful (D).
(3) In round 3, each server decides whether to report collusion, R or $\bar{R}$.

Solution concept. The game we have defined is a sequential game of complete information, i.e., players know others' past actions. Nash equilibria, the standard notion of equilibrium in game theory, are known to predict "unreasonable" behaviors in sequential games. In a Nash equilibrium of a sequential game, a player can commit to empty/non-credible threats.[2] We then consider a refinement of Nash equilibrium for sequential games, namely *Subgame Perfect Equilibrium* [80] (SPE). In SPE, a player chooses the optimal action at any point, i.e., the action that maximizes expected utility from that point on, no matter what has happened in the past.

Earning privacy worth. Let V be the worth of the user secret. If the game stops in round one, each queried server only receives service fee λ_s. Otherwise, if they all play amicable (A) in round two, they all learn the secret and receive V. If at least two parties play D, then they do not learn the secret. If there is exactly one deceitful party, then others do not learn the secret and the deceitful server learns the secret in its best case, e.g., when colluders exchange queries or responses, but not otherwise.

SPE. In solving the collusion game G, we let $\lambda_r = 0$. The result for the single-shot k-party collusion game is as follows.

[2]Consider the Ultimatum game with two players. One player proposes a way to divide a sum of money between them. If the responding player accepts the proposal, then the money is divided accordingly; otherwise, both receive nothing. The responder can threaten to accept only fair offers, but this threat is not credible: given any non-zero amount, the only rational action for the responder is to accept.

Theorem III.1. *In the k-party collusion game* G *in* $(\ell, k, 1)$-*PIR, the strategy profiles of all ℓ servers playing $\bar{C}D$ or CD in rounds one and two are the only two SPEs when $\lambda_p > 0$ and $\lambda_s + \frac{k-1}{k}\lambda_p > V$.*

In Theorem III.1, we only describe the strategies in the first two rounds that are pertinent to our desired non-collusion outcome, where servers do not collude or input gibberish into collusion protocols.

Proof. First, consider $k = 2$. The extensive form game tree for 2-server collusion is shown in Figure 2. We depict the good case for the only deceitful player where it learns the secret in Figure 2 and consider *both good and bad cases* where it does not learn the secret, in analysis. The bad case can occur if there is a conditional abort upon detection of incorrect inputs.

At decision node 12.3 and 12.4, it is straightforward to determine $\alpha_4 = 1$ and $\beta_4 = 1$ since deviating does not generate more returns when $\lambda_r = 0$. This means that the deceitful server reports an amicable colluder. At decision node 12.2, to make R rational, we need R to generate higher expected returns:

$$R : \beta_3(V - \frac{\lambda_p - \lambda_r}{2}) + (1 - \beta_3)(V + \lambda_s + \lambda_r)$$
$$= V + \lambda_s - \beta_2(\frac{\lambda_p}{2} + \lambda_s)$$
$$\bar{R} : \beta_3(V - \lambda_p) + (1 - \beta_3)(V + \lambda_s)$$
$$= V + \lambda_s - \beta_2(\lambda_p + \lambda_s)$$

Playing R produces higher expected returns since $\lambda_p > 0$. Then $\alpha_3 = \beta_3 = 1$. Working backward, at decision node 12.1, we consider the following:

$$\text{Good case: } A : \beta_2(V - \frac{\lambda_p - \lambda_r}{2}) + (1 - \beta_2)(-\lambda_p)$$
$$= -(1 - \frac{\beta_2}{2})\lambda_p + \beta_2 V$$
$$D : \beta_2(V + \lambda_s + \lambda_r) + (1 - \beta_2)\lambda_s = \lambda_s + \beta_2 V$$
$$\text{Bad case: } A : \beta_2(V - \frac{\lambda_p - \lambda_r}{2}) + (1 - \beta_2)\lambda_s$$
$$= \beta_2(V - \frac{\lambda_p}{2} - \lambda_s) + \lambda_s$$
$$D : \beta_2\lambda_s + (1 - \beta_2)\lambda_s = \lambda_s$$

Playing D generates higher returns in good case and at least the same returns in bad case since $\lambda_s + \frac{\lambda_p}{2} > V$. Then $\alpha_2 = 0$ and similarly, $\beta_2 = 0$. At decision node 12.0, playing $\bar{C}$ gives the same returns as playing C, i.e., service fee λ_s. Then $\alpha_1 = \beta_1 = 0$ or $\alpha_1 = \beta_1 = 1$. Hence, the SPEs for $k = 2$ are $\bar{C}DR\bar{R}_3R_4$ and $CDR\bar{R}_3R_4$.

Next, consider $k > 2$. The example extensive form game tree for $k = 3$ can be found in Figure 3. In round three, it is straightforward to see that when there is more than one deceitful player, all parties can only play $\bar{R}$. When there is exactly one deceitful player, e.g., decision node 123.3 in Figure 3, this player plays R and others play $\bar{R}$. When all players are amicable, e.g., node 123.2, the expected returns from playing R are greater than playing $\bar{R}$ when $\lambda_p > 0$ because $V - \frac{k-1}{k}\lambda_p >$
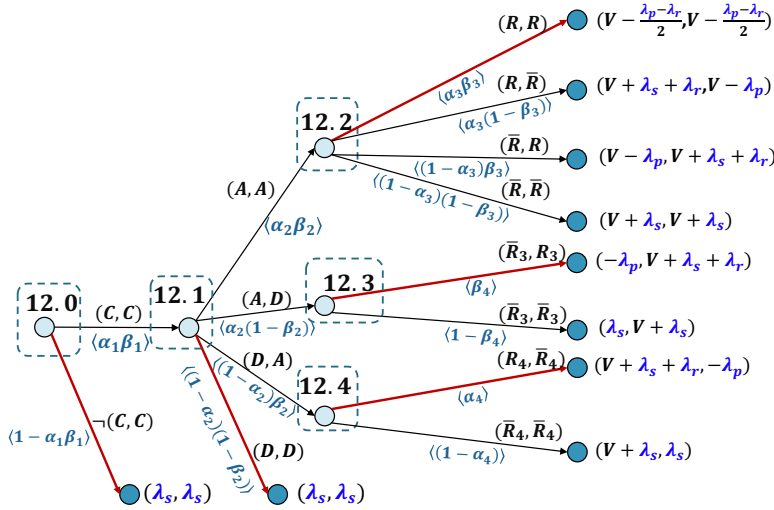
5

Fig. 2: Extensive form of G with $k = 2$. Denote the queried servers as S_1, S_2. Each node enveloped with dashed lines is a decision node, with the prefix of its name indicating the owner (e.g., "12" means that S_1 and S_2 make decisions at this node). Actions of servers are specified on the one-directional arrows and the probability for each move is under the arrows in brackets, with αs for S_1, and βs for S_2. We label the round three actions R and $\bar{R}$ at decision nodes 12.3 and 12.4 respectively as $R_3, \bar{R}_3, R_4, \bar{R}_4$. The payoff vector is in the form (x, y) with x for S_1 and y for S_2. When both servers report, we apply an equal probability of submitting first. The thick red edges indicate the SPE.
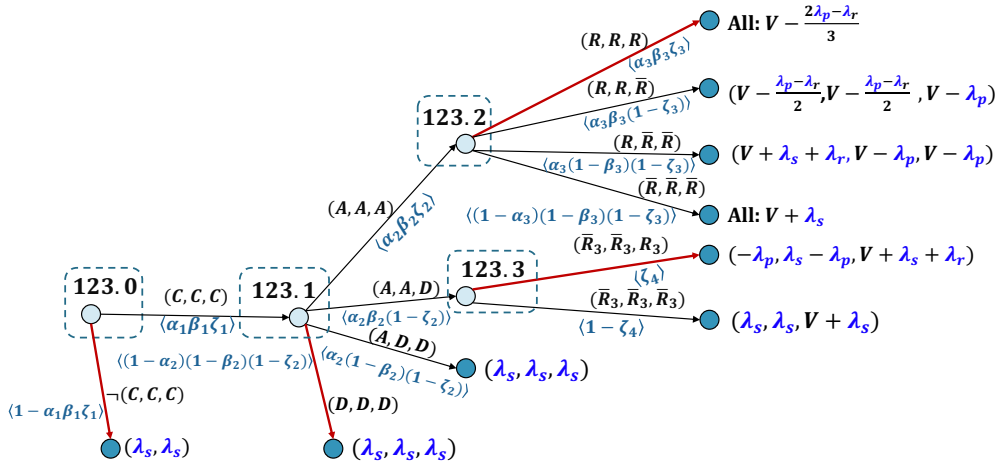

Fig. 3: Extensive form of G with $k = 3$. We label the round three actions R and $\bar{R}$ at decision nodes 123.3 as $R_3, \bar{R}_3$. The thick red edges indicate the SPE.

$V - \lambda_p$, similar to the $k = 2$ case. Next, in round two, let x_1 be the probability of exactly one other party playing D ($x_1 \in [0, 1]$), and x_2 be the probability of at least two other parties playing D ($x_2 \in [0, 1 - x_1]$). We can compute the returns from playing A and D as follows:

Good case: $A : x_1(-\lambda_p) + x_2\lambda_s + (1 - x_1 - x_2)\cdot$
$$(V - \frac{(k-1)\lambda_p - \lambda_r}{k})$$
$D : x_1\lambda_s + x_2\lambda_s + (1 - x_1 - x_2)(V + \lambda_s + \lambda_r)$

Bad case: $A : x_1\lambda_s + x_2\lambda_s + (1 - x_1 - x_2)\cdot$
$$(V - \frac{(k-1)\lambda_p - \lambda_r}{k})$$

$$D : x_1\lambda_s + x_2\lambda_s + (1 - x_1 - x_2)\lambda_s$$

It is straightforward to see that playing D produces higher expected returns. Then in round one, playing $\bar{C}$ produces the same returns as playing C. Therefore, playing $\bar{C}D$ or CD in rounds one and two are the SPEs. $\square$

Counter false accusations. (1) For random accusations, for now, we focus on $f(\cdot)$ that reveals explicit bits about user secrets and consider a more general case in Section VI-B. A server can then make successful random accusations with probability $\frac{1}{2}$. (2) An informed queried server knows some non-trivial function about the user secret. It can falsely accuse others successfully with probability $\frac{\omega+2}{2(\omega+1)}$ ($= \frac{1}{\omega+1}\cdot 1 + \frac{\omega}{\omega+1}\cdot\frac{1}{2}$). Com-

bining the two cases, we require $(1 - \max\{\frac{1}{2}, \frac{\omega+2}{2(\omega+1)}\})\lambda_f > \max\{\frac{1}{2}, \frac{\omega+2}{2(\omega+1)}\}\lambda_r$ or simply $\lambda_f > 0$ so that servers do not falsely accuse others. (3) A user can also collude with a queried server to make fake accusations. By having $(k-1)\lambda_s > \lambda_r = 0$ or simply $\lambda_s > 0$, the user is discouraged from falsely accusing servers.

Comments on the simultaneous moves. In round two, the parties select action A or D simultaneously. Alternatively, if simultaneity cannot be implemented, one party has to move first. Then no matter what action this initiator player takes, the follower's dominant strategy is D. The SPEs stay the same. We give more advantages to colluders by allowing simultaneous moves.

Blind collusion. When parties play D, without embedded verification gadgets, it is computationally hard to fabricate seemingly valid inputs for many PIR protocols, e.g., the Boyle-Gilboa-Ishai and Hafiz-Henry protocols we consider, while being (non-negligibly) possible for others. We focus on collusion with verification and the computationally hard case without verification in the main body. Appendix A addresses the latter blind collusion without verifying inputs. There, we solve for the conditions on parameters to achieve the non-colluding sequential equilibrium [58] in Theorem A.1.

B. Non-collusion outcome in repeated games

Consider the following repeated collusion game $\mathsf{G_R}$ for a player S_i. In each run $1, \ldots, T$:

 (0) In round zero, Nature decides whether S_i is queried. If not, S_i terminates this run and enters round zero of the next run. Otherwise, S_i enters game G.

When T is finite and known, we can solve for the equilibrium using backward induction as before. In the last run of the game where S_i is queried, regardless of history, the SPEs are the same as before under the same set of conditions, i.e., S_i reports after successful collusion in round three, plays D in round two, and plays C or $\bar{C}$ in round one. With this knowledge, players do not deviate from the prescribed SPEs in the single-shot game in the second-to-last run, and so on, until the first run. We state the following corollary and omit the proof since it follows from Theorem III.1.

Corollary 1. *The ℓ-party repeated collusion game $\mathsf{G_R}$ in $(\ell, k, 1)$-PIR with a known finite number of runs has the same SPEs under the same conditions as in Theorem III.1.*

When T is infinite or finite but unknown, backward induction no longer applies and collusion can become an equilibrium [64]. For simplicity, we address both cases as infinitely repeated games. The Folk Theorem [37] states that if players are patient enough, then repeated interaction can result in virtually any average payoff in an SPE. If the players are memoryless (i.e., always play like they are playing for the first time), we can still achieve the non-colluding equilibrium. We show the same result *without* making such additional assumptions. Specifically, we consider player strategies of the form "If another player has reported collusion, never collude with this player again". Because strategies of this form yield undesirable equilibria in

many repeated versions of well-studied games, e.g., repeated prisoner's dilemma.

SPE. In solving game $\mathsf{G_R}$ with unknown or infinite T, we let the reward $\lambda_r > 0$. Same as before and after simplification, we require $\lambda_f > \frac{\omega+2}{\omega}\lambda_r$ and $\lambda_s > \frac{1}{k-1}\lambda_r$ so that servers and users do not falsely accuse others. We now formally state the main result for the ℓ-party repeated collusion game. We apply discount rate $\delta \in [0, 1)$ on future returns [71, Chap. 7.2] because rational agents prefer acquiring the same returns earlier than later. Higher δ indicates higher patience.

Theorem III.2. *In the ℓ-party infinitely repeated collusion game $\mathsf{G_R}$ in $(\ell, k, 1)$-PIR, the strategy profile of all ℓ servers playing $\bar{C}D$ or CD in rounds one and two of each run are the only SPEs when $\frac{\delta}{1-\delta}(1-q)V < \lambda_r \leq \lambda_p$, and $\lambda_s + \lambda_p > \frac{1}{k}(\lambda_r + \lambda_p) + V$ where $q = \frac{(\ell-k)\cdots(\ell-2k+2)}{(\ell-1)\cdots(\ell-k+1)}$.*

We only describe the strategies in the first two rounds relevant to the non-collusion outcome. With the same k, when we increase ℓ, we increase q, which expands the feasibility region for λ_r.

Proof. When playing $\mathsf{G_R}$, suppose for contradiction servers play $CA\bar{R}$ in equilibrium, i.e., collude without reporting, and they always refuse to collude with those who have reported their collusion in the past. Consider a server S_1 that deviates to play R in round three of the collusion game G in a run. Let p_i be the probability that server S_1 can collude with the other $(k-1)$ queried servers in the i-th run where S_1 is queried $(i \geq 1)$, meaning that none of them have reported each other before.

First, consider $k = 2$. Two queried servers collude only if they have never met before. In the i-th run, suppose S_1 and S_2 are queried. The probability that in one of the previous $(i-1)$ instance, P_2 is not queried is $\frac{\ell-2}{\ell-1}$. Then for $k = 2$, $\mathsf{p}_i = (\frac{\ell-2}{\ell-1})^{i-1}$ since the runs are independent when each client samples servers at random by themselves.

Next, consider $k > 2$. Unlike in the $k = 2$ case, the $k - 1$ queried servers collude only if they have never colluded with S_1 before, even though they may have already met in some runs. This is because collusion may have not happened due to some other queried servers refusing to collude with S_1. We can obtain a lower bound on p_i by considering collusion only if the k parties have never met S_1 before. This gives $\mathsf{p}_i \geq \left[\frac{\binom{\ell-k}{k-1}}{\binom{\ell-1}{k-1}}\right]^{i-1} = \left[\frac{(\ell-k)\cdots(\ell-2k+2)}{(\ell-1)\cdots(\ell-k+1)}\right]^{i-1}$. Let $q = \frac{(\ell-k)\cdots(\ell-2k+2)}{(\ell-1)\cdots(\ell-k+1)}$. Then $\mathsf{p}_i \geq q^{i-1}$.

From playing $CA\bar{R}$, S_1 earns service fees and the privacy worth in each run:

$$\lambda_s + V + (\lambda_s + V)\sum_{i=1}^{\infty}\delta^i$$

From reporting collusion, S_1 expects to earn rewards in the runs where colluding partners are strangers and at least the service fees otherwise:

$$V + \lambda_s + \lambda_r + \sum_{i=1}^{\infty}(\mathsf{p}_{i+1}(V + \lambda_s + \lambda_r) + (1 - \mathsf{p}_{i+1})\lambda_s)\delta^i$$

Subtracting the first quantity from the second gives

$$\lambda_r + (V + \lambda_r) \sum_{i=1}^{\infty} \mathsf{p}_{i+1}\delta^i - V\frac{\delta}{1-\delta} \geq$$

$$\lambda_r + (V + \lambda_r) \sum_{i=1}^{\infty} (q\delta)^i - V\frac{\delta}{1-\delta}$$

Then, as long as $\frac{1}{1-q\delta}\lambda_r > (\frac{\delta}{1-\delta} - \frac{q\delta}{1-q\delta})V$, playing $\bar{R}$ cannot be equilibrium strategy since playing R is a profitable deviation. This is satisfied when $\lambda_r > \frac{\delta}{1-\delta}(1-q)V$. This means that one reports each successful collusion. Then following from Theorem III.1, in round two of each collusion game, playing D generates higher returns, and in round one, playing $\bar{C}$ or C generates the same returns. Playing $\bar{C}D$ or CD in the first two rounds of each run are the only SPEs. □

What is special about infinite repetition. The theory of infinitely repeated games often faces a multiplicity of equilibria [26] dependent on players' patience, i.e., how much they discount future returns. For example, in a single or finitely repeated two-person prisoners' dilemma, both players playing "defect" is the Nash equilibrium; in an infinitely repeated prisoners' dilemma, "cooperate" becomes the equilibrium strategy when they are sufficiently patient and discount future returns by a sufficiently small amount. This is because there is no end game, and a player can always penalize "defect" actions. Except for penalties, in broader social contexts, equilibrium different from those appearing in finite games can also arise due to reputation concerns in long-run relationships and shared expectations of social norms. Examples include neighbors in Shasta County, CA, resolving disputes caused by each other's cattle via informal social norms instead of legal rules without any hierarchical coordinator [33], and anonymous sellers of a certain size in Deep Web exhibiting honest patterns to amass good reputations without government regulations [48].

Our goal is then to preserve the equilibrium in finite games in infinite runs, even when players are highly patient. The first step is to obstruct the formation of informal institutions [49] or long-term relationships among the players by having many servers and querying only some of them in each run. Combining this with positive rewards, we increase players' returns from following the prescribed equilibrium strategy in the current run and decrease the probability of them playing together again. In this way, even a patient player expects higher utility from observing the desired non-colluding strategy.

C. Existence of solution

Now one natural question is whether practical solutions are guaranteed to exist. To ensure affordable service fees λ_s, we define practical service fees mathematically as $\lambda_s \leq \frac{V}{k}\xi$, where $\xi \in (0,1)$ is a practicality parameter characterizing the affordability of λ_s. We state the following proposition.

Proposition 1 (Existence of solution). *For $\ell \gg k$, there always exist parameter assignments for inequalities ①-⑤: ① $\lambda_f > \frac{\omega+2}{\omega}\lambda_r$; ② $(k-1)\lambda_s > \lambda_r$; ③ $\frac{\delta}{1-\delta}(1 - \frac{(\ell-k)\cdots(\ell-2k+2)}{(\ell-1)\cdots(\ell-k+1)})V <$*
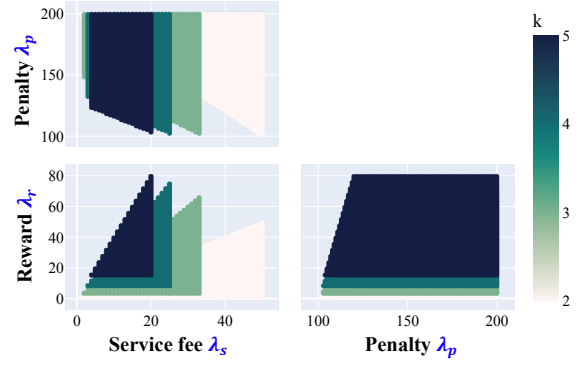

Fig. 4: Example parameter feasibility regions. We provide a 3D interactive plot here [42].

$\lambda_r \leq \lambda_p$; ④ $\lambda_s + \lambda_p > \frac{1}{k}(\lambda_r + \lambda_p) + V$; ⑤ $\lambda_s \leq \frac{V}{k}\xi$, where $\delta \in [0,1), \xi \in (0,1)$.

Proof. To satisfy inequality ② and ⑤, we need $(k-1)\frac{V}{k}\xi > \lambda_r$, which gives $\xi \geq \frac{k\lambda_r}{(k-1)V}$. We can have a feasible $\xi < 1$ by mandating $k\lambda_r < (k-1)V$. Considering ③, we have feasible assignments for the reward amount if $\frac{\delta}{1-\delta}(1-(\frac{\ell-2k+2}{\ell-k+1})^{k-1}) < \frac{k-1}{k}$. Then we can pick λ_f and λ_p according to inequality ①, ③ and ④, λ_s according to inequality ②, ④ and ⑤. □

An example feasible parameter region of λ_s and λ_p in setting $\ell = 10000, V = 100, k = \{2,3,4,5\}, w = 1, \lambda_p \leq 2V, \delta = 0.99$ is shown in Figure 4. Higher service fee allows for a lower penalty (④) and higher reporting reward (②). Lower δ will expand the feasibility regions (③) because players are less patient and prefer receiving the same payment earlier. One sample solution here is $k = 2, \lambda_s = 1, \lambda_f = \lambda_p = 200, \lambda_r = 0.99$.

IV. MALICIOUS SERVERS

We now consider $\theta \in [0,1)$ of the ℓ servers to be malicious. Malicious servers behave arbitrarily and ignore incentives. They bring about two changes to the mechanism. First, the incentive mechanism achieves the non-collusion outcome only probabilistically. Second, the collusion evidence verification routines described in Figure 1 can be exploited by malicious servers to function as an expensive oracle for revealing user secrets.

A. Statistical non-collusion outcome

Consider a statistical security parameter η, e.g., $\eta = 40$. To achieve the non-collusion outcome with probability at least $1 - 2^{-\eta}$, θ needs to satisfy

$$\frac{\binom{(1-\theta)\ell}{0}\binom{\theta\ell}{k}}{\binom{\ell}{k}} + \frac{\binom{(1-\theta)\ell}{1}\binom{\theta\ell}{k-1}}{\binom{\ell}{k}} \leq 2^{-\eta} \qquad (S^\star)$$

so that at least two of the queried servers are rational. Following from Theorem III.1 and III.2, we have the following corollaries.

Corollary 2. *In the k-party collusion game* G *in* $(\ell, k, 1)$-*PIR with* $\theta\ell$ *malicious servers, with probability* $1 - 2^{-\eta}$, *the strategy profile of all* $(1 - \theta)\ell$ *rational servers playing* $\bar{C}D$ *or* CD *in rounds one and two are the only SPEs when* $\lambda_p > 0$ *and* $\lambda_s + \frac{1}{2}\lambda_p > V$ *where* θ *satisfies* $S^\star$.

Corollary 3. *In the* ℓ-*party repeated collusion game* G_R *with an infinite or unknown number of runs in* $(\ell, k, 1)$-*PIR with* $\theta\ell$ *malicious servers, with probability* $1 - 2^{-\eta}$, *the strategy profile of all* $(1 - \theta)\ell$ *rational servers playing* $\bar{C}D$ *or* CD *in rounds one and two of each run are the only SPEs when* $\frac{\delta}{1-\delta}(1 - q)V < \lambda_r \leq \lambda_p$, *and* $\lambda_s + \lambda_p > \frac{1}{2}(\lambda_r + \lambda_p) + V$ *where* $q = \frac{(\ell-k)\cdots(\ell-2k+2)}{(\ell-1)\cdots(\ell-k+1)}$ *and* θ *satisfies* $S^\star$.

B. Privacy-preserving verification

The second change is that now we need to protect user secrets during collusion resolution. This is not a concern when servers are rational since they do not collude in equilibrium. But malicious parties do not respond to incentives, they can falsely report collusion and learn the secret by paying the penalty λ_p. Therefore, in collusion verification, instead of having involved parties reveal commitments and perform verification on plaintexts, we adopt zero-knowledge accusation verification.

Upon receiving an accusation with a nontrivial $f(\cdot)$, **CC** counts down a proof collection window. The accused parties collectively provide a zero-knowledge succinct non-interactive argument of knowledge (zkSNARK) π justifying the evidence being false (line 2 in Figure 5). **CC** confirms the accusation and executes payments if the accused servers fail to provide proof in time Δ^* or if the proof is rejected (line 3). We discuss setting Δ^* in Section VI-B. If the accuser somehow composes a zkSNARK π^* for the knowledge of inputs and outputs of the function $f(\cdot)$ during the executions of the collusion protocol, π^* can be submitted as evidence and the contract checks π^* directly.

Zero-knowledge Collusion Resolution (S_{id_i} accusing)

1) Server S_{id_i} submits a collusion report to **CC**, indicating the accused parties and evidence, i.e., circuits for computing $f(\cdot)$, S_{id_i}'s inputs to $f(\cdot)$, and the commitment to the function outputs. **CC** rejects the report if $f(\cdot)$ is trivial.
2) Accused server(s) provide π proving that the committed value in the evidence does not match the function output of $f(\cdot)$.
3) **CC** confirms the accusation if (1) the accused server(s) fails to submit the proof in time Δ^*, or (2) π is rejected.

Fig. 5: Overview of zero-knowledge collusion resolution.

Constructing zkSNARKs. When the evidence in the accusation is a query or response that has been committed to on the board **BB**, the server being accused can submit *proof of inequality*. If the accuser does not specify the accused, all other servers queried in this instance need to provide inequality proofs. More concretely, suppose we work with Pedersen commitment with common generators g, h and modulus N. It is perfectly hiding and satisfies our needs. Let $c = g^m h^r$ be a commitment sent

by a party S_x on **BB** where m is the committed message and r is the randomness. A party accuses with evidence (g^s, S_x), i.e., "the exponent of g^s is the committed value in c", instead of revealing s directly. Then S_x needs to prove the statement that $c' = c/g^s$ is not a commitment of zero, i.e., proof of knowledge $PK\{(a, r) : c' = g^a h^r \wedge a \neq 0 \wedge r \neq 0\}$ where $a = m - s$. This can be efficiently proven with range proofs [14].

Next, consider the scenario where the evidence is a general function circuit along with its output. There are two scenarios. First, if $k = 2$, the accused server generates the proof on its own. Suppose we still work with Pedersen commitment with parameters g, h, N. Now suppose a party accuses with general evidence stating that "the exponent of g^s is the output of $f(\cdot)$" along with its collusion input m_1, which is committed on **BB** in a commitment c_1. Let the accused server hold the corresponding commitment $c_2 = g^{m_2} h^r$. She proves $PK\{(m_2, r) : c_2 = g^{m_2} h^r \wedge g^a = g^{f(m_1, m_2)} \wedge g^a \neq g^s\}$. For simplicity, we let $f(\cdot)$ be expressed as an arithmetic circuit since the accused party can arithmetize it if $f(\cdot)$ is a Boolean circuit. We can then employ applicable schemes including Groth16 [45], Plonk [38], or LegoSNARK [16] with a polynomial commitment. Second, when $k > 2$, the accused servers can utilize the collaborative zkSNARKs primitive [74] to generate the proof while preserving privacy for their inputs. The knowledge to prove is similar to before except that the inputs to $f(\cdot)$ are held by multiple parties.

V. EXITING STRATEGIES AND COALITION FORMATION

In this section, we analyze how servers' exiting strategies and strong coalition formation affect the proposed mechanism and its parameterization.

A. Exiting strategy

We first study the two exiting strategies, carrying out massive collusion before exiting and performing collusion or data selling after leaving the system.

A server S_e is leaving the system. The maximum fines we can take from S_e is the deposited penalty λ_p. We propose a Self-Insurance [32] design tailored to our needs and detain service fees as the self-insurance pool for a certain privacy protection period of T time units. Suppose during the period, Ω users submit $k\Omega$ total queries. The accumulated service fees are $k\Omega\lambda_s$. The expected number of queries that S_e receives is $\frac{k\Omega}{\ell}$. Then S_e can recover $\frac{k\Omega}{\ell}$ user secrets in expectation in its best case. Consider the worst case (for us) where all collusion is successful. We expect $(k - 1)/k$ of the collusion to be reported by colluding servers other than S_e. The expected total amount of fines we cannot realize from S_e is $(\frac{k\Omega}{l}\frac{k-1}{k} - 1)\lambda_p = \frac{(k-1)\Omega - \ell}{\ell}\lambda_p$. In the meantime, the service fee pool also grows with Ω. There are $k \cdot \frac{k\Omega}{\ell}\lambda_s$ transaction fees that are available to cover the potential loss from S_e's exiting. We use

$$\sigma = \frac{\frac{(k-1)\Omega - \ell}{\ell}\lambda_p}{k\frac{k\Omega}{\ell}\lambda_s} = \frac{(k-1)\Omega - \ell}{k^2\Omega}\frac{\lambda_p}{\lambda_s}$$

to denote the tension between the size of inexecutable fines and the service fee pool, and $\sigma \leq 1$ is needed to cover fines

with detained service fees. ℓ being significantly larger than k can allow a small fee amount λ_s.

To be more practical, we can assume an interest rate $r \in [0, 1]$ per time unit on these accumulated service fees, which then become $k\Omega\lambda_s(1+r)^T$. We can set the privacy protection period to be T and discount the penalty across time by rate $r' \in [0, 1]$, which means fewer fines as a query becomes ancient for a positive r'. We then obtain

$$\sigma' = \frac{(k-1)\Omega - \ell}{k^2\Omega} \frac{\lambda_p(1-r')^T}{\lambda_s(1+r)^T}$$

where $\sigma' \leq 1$ is also desired to relieve the tension. As an illustration, let $k = 2, \ell = 1000, T = 10000, \Omega = 5000$. We have $\sigma' = \frac{1}{5}\frac{\lambda_p(1-r')^{10000}}{\lambda_s(1+r)^{10000}}$. For $r' = r = 0.0001$, we can have $\sigma' \leq 1$ when $\lambda_s \geq 0.027\lambda_p$. Note that this is not the exact formulation as we are treating the queries as coming at once at the beginning of the T time units. We provide a detailed derivation and consider more servers exiting in Section C.

A server has left the system. After a server has exited the service, it can sell gathered information to parties inside or outside the system without being penalized. This is outside the scope of this work, but we can enforce the rule that servers only receive their deposits back after the privacy protection period ends for their last received query. Additionally, the proof of secure erasure (PoSE) technique [76] that ensures the erasure of memory can potentially counter this threat without keeping servers waiting. We can require a PoSE attestation upon a deposit withdrawal request.

B. Strong coalitions

We now extend the analysis to consider strong coalitions formed by servers. The threat model for group players is similar to the individual server case but has three major changes. (1) Servers inside the coalition act in a highly coordinated way. (2) Information entering into the group is shared among members, and (3) members do not report collusion inside.

1) Bound the loss within the current model: Suppose the pre-established coalition is of size s, and that the number of servers to be queried by the user is k ($\leq \ell$). We consider the following two cases. First, if $s < k - 1$, at least two external rational servers are queried. Collusion with 2 rational outsiders is discouraged by the current mechanism.

Second, if $s \geq k - 1$, with some probability p_s, the clique receives at least $k - 1$ of the queries. We can calculate $p_s = \frac{\binom{s}{k-1}\binom{\ell-s}{1}+\binom{s}{k}}{\binom{\ell}{k}}$. Then the expected loss from privacy breach when $s \geq k - 1$ is p_sV. Note that we do not explicitly include malicious servers in coalitions. But they can be incorporated by conceptually adding $\theta\ell$ more members into the coalition.

2) Update the model: Cooperative game theory [12] studies the formation of coalitions. Our goal is to cap the coalition size. A cooperative game can be formalized with the set of ℓ players $S = \{S_i | i \in [\ell]\}$, each nonempty set of coalitions, and a characteristic / value function v mapping a coalition to its collective payoff, $v : 2^S \mapsto \mathbb{R}$. For example, the grand coalition has value $v(S) = \ell V + k\lambda_s$. An allocation $\mathsf{x} \in \mathbb{R}^\ell$ for a coalition specifies how to divide the coalition value among members.

Solution concept. Suppose the coalition structure (partition of S) is non-overlapping. We solve for the allocation in the core of any coalition C of size s, denoted as $v(C)$. A core allocation is a feasible allocation that no other sub-coalition ($\subset C$) can improve on (i.e., give a strictly higher payoff for all). Here, feasibility simply requires that one does not over-allocate the coalition value. For any C with size $s \geq k$, its value function is $(p_s \cdot sV + k/\ell \cdot s\lambda_s)$. Then it is straightforward to see that its core is allocating the coalition value evenly: $\{\forall S_i \in C, \mathsf{x}(S_i) = p_sV + k/\ell \cdot \lambda_s\}$. This leads to the grand coalition being the most desired.

Alternatively, consider a general value function where the coalition C earns $v(C) = p_sD(s) + k/\ell \cdot s\lambda_s$ where $D(s)$ is the total value earned by group C. We let $D(s)$ be non-increasing in s. This is reasonable when the user secrets are rivalrous where a party values secrets more when fewer parties also learn, e.g., upcoming valuation changes of a company in the stock market and hidden treasures. In such scenarios, individual profits that can be capitalized decrease as more parties learn this information.

First, consider constant coalition value $D(s) = cV$ where $c > 0$ is a constant. In this setting, the core is $\{\forall S_j \in C, \mathsf{x}(S_j) = p_s/s \cdot cV + k/\ell \cdot \lambda_s\}$. By including one more member to C, one existing member's extra gain is $(p_{s+1}/(s+1) - p_s/s) \cdot cV$. It is straightforward to see that the quantity is always negative when $k = 2$. This indicates that the coalition does not grow in size for $k = 2$. By setting $k = 2$, we can limit coalition size to 2. For $k \geq 3$, by making the extra returns negative, we can obtain the maximum coalition size $\lfloor\frac{(k-2)(k\ell-k+1)}{(k-1)^2}\rfloor$. This means that the coalition is incentivized to include new members until it reaches this maximum size.

Similarly, consider linear depreciation of coalition value with respect to coalition size, i.e., $D(s) = cV/s$. The core is $\{\forall S_j \in C, \mathsf{x}(S_j) = p_s/s^2 \cdot cV + k/\ell \cdot \lambda_s\}$. Adding one more member to C increases p_s while reducing privacy gain. An existing member's extra profit from including a new member is $(p_{s+1}/(s+1)^2 - p_s/s^2) \cdot cV$. For $k = 2$, this quantity is negative, and the coalition does not have an incentive to include new members. By setting $k = 2$, we limit coalition size to 2. For $k > 2$, we can bound the size of the coalition by enforcing negative extra profits as before. Let $c_1 = -(k-1)(k-2), c_2 = (\ell - 2)k^2 - (3\ell-7)k - 5, c_3 = (k-2)(k\ell-k+1), \Delta = c_2^2 - 4c_1c_3$, then the coalition includes new members until its size reaches $\lfloor\frac{-c_2 - \sqrt{\Delta}}{2c_1}\rfloor$. We provide a simple graphical demonstration for locating the upper bound of coalition size here [43]. One can continue to consider a higher value depreciation rate, which allows for further limiting coalition sizes for $k > 3$.

VI. GAME DESIGN FLOW AND IMPLEMENTATION

A system designer first evaluates the participating servers and determines a recommended k. She then decides the privacy value upper bound V to protect. One can compute the parameters for the collusion game according to inequalities in Proposition 1. It outputs a list of feasible value assignments

for parameters. There is no absolute standard to select one assignment over another but rather depends on the priorities.

For illustration purposes, we present abstract essential functions for implementing the game and payment rules in Algorithm 1. In Line 15, a user submits queries to a list of servers. The coordinator initializes parameters for future possible accusation resolutions. In Lines 21-24, an accusation is filed, and we call the accusation validation routine. Line 4 checks the validity of the accusation. Line 5 calls the payment execution routine to realize payment rules.

ALGORITHM 1: Coordinator (Essential Functions)

Input: $\lambda_r, \lambda_p, \lambda_s, \lambda_f$

1 **Function** AccusationVal ($witness, id, \text{ed}$) **:**
2   **for** $pk \in$ Journal$[id].pkList$ **do**
3   **if** !alreadyFined$[id][pk]$ *and* ed *is not recorded in* Journal$[id].filed[pk]$ **then**
4   **if** ValidationUtilityFunc (pk, id, ed) **then**
5   ▷ Record revealed valid evidence
6   Journal$[id].filed.append(witness, \text{ed})$
7   PaymentExec ($witness, id$)
8   **else**
9   Take fines λ_f from Journal$[id].witness$
10 **Function** PaymentExec ($witness, id$) **:**
11   Read from Journal$[id]$ the request $sender, pkList$
12   **for** $pk \in pkList \ \wedge \ pk != witness$ **do**
13   Take penalty λ_p from pk
14   alreadyFined$[id][pk] = True$
15   Distribute reward λ_r to $witness$
16 **On new** *incoming message* m**:**
17   **if** m *is a new query* **then**
18   Lock λ_s from m.$sender$ for each $pk \in$ m.$pkList$
19   Issue a unique identifier id for the request
20   Journal$[id] \leftarrow$ (m.$sender$, m.$requests$, m.$pkList$)
21   **for** $pk \in$ m.$pkList$ **do**
22   alreadyFined$[id][pk] = False$
23   **if** m *is an accusation* **then**
24   Lock λ_f from accuser m.$sender$
25   AccusationVal (m.$sender$, m.id, m.ed)

A. Implementation

Smart contract implementation of CC. We treat Ethereum as a public bulletin board and conveniently implement the coordinator contract as a smart contract on Ethereum. The source code in Solidity is made available [44]. The contract maintains the complete life cycle of PIR service and resolves collusion accusations. For evidence verification, as shown in Figure 1, the accuser submits the evidence, followed by the involved server(s) submitting auxiliary information. If no supplementary data is submitted in time, the accusation is automatically marked as successful. The accused server can be pinpointed through elimination in this case. After collecting information, the verification algorithm computes certain function circuits after verifying function non-triviality and input validity. Any secure commitment and signature schemes suffice, e.g., cryptographic hash functions like SHA-3 as commitment scheme and ECDSA signatures [55].

To facilitate evidence verification for a general $f(\cdot)$, we provide a sample reconstruction function as a default $f(\cdot)$ inside the contract. So the accuser can omit the step of providing a function circuit if $f(\cdot)$ is the reconstruction algorithm. Otherwise, we ask the accuser to address the function in a predetermined way, e.g., $f(bytes\ memory[k])$, when providing the circuit for a simpler calling convention. We summarize the gas costs of contract deployment and each function call in Table I in Section VI-A.

The gas costs of one-time contract deployment and function calls are summarized in Table I, which we find to be acceptable, e.g., posting requests cost about 0.74 dollars [85]. Here, CheckCircuits$(\cdot)$ examines the non-triviality of functions mentioned in accusations, and "+" represents payments to oracle services (discussed below). VerifyExchange$(\cdot)$ verifies collusion when the evidence is a de-committed query or response. VerifyGeneralFunc$(\cdot)$ verifies collusion for a general $f(\cdot)$. zkVerify$(\cdot)$ performs zero-knowledge verification when malicious servers exist (discussed in Section IV). The costs induced by the base PIR service, i.e., fair exchange of service and fees, are approximately the "Normal service" functions. The costs added by the mechanism are transmitting fixed-length commitments of the random companion queries and the functions for "Collusion resolution".

Cryptographic, communication and computation overhead. We preserve the communication complexity of the underlying multi-server PIR because participants only send an additional commitment and signature for each message. Computing SHA-3 hash costs 30 base gas and 6 gas every 256 bits [34], and verifying ECDSA signatures involves recovering the signer with ecrecover, which consumes 3000 gas. In terms of computation overhead, SHA-3 is efficient, and an ECDSA signature can be computed in about 0.1ms [24]. This delay in generating responses is not significant. In two efficient single-server cPIR solutions, XPIR [1] and SealPIR [5], it takes a server three orders of magnitude more time to compute the response even for small databases ($N = 2^{16}$).

Note that the gas costs related to verifying signatures and messages are not induced by collusion mitigation, but to counter malicious behaviors of sending incorrect messages in general. Existing methods for combating the problem of malicious server sending incorrect messages brings even higher inefficiency: naively applying GMW compiler [41] invokes higher communication costs than sending the database; robust multi-server solutions [8], [27], [40] are not generally applicable to any existing schemes, do not support database update, and also have high communication complexity; verifiable PIR [95] relies on cryptographic assumptions and requires a trusted third party to generate verification keys.

Non-triviality of function circuits. When parties report with evidence for a general $f(\cdot)$, we examine the non-triviality of the function $f(\cdot)$ in addition to verifying function inputs and outputs. We view a function as non-trivial if its output depends on inputs, and taint analysis [73] can be employed for this purpose. When we implement the design on Ethereum with smart contracts, the functions do not contain pointers and are deterministic.

However, accomplishing this task on-chain can be impracti-

TABLE I: Cost estimates in Gas and dollars [85].

Normal service	Gas	Dollars	Collusion resolution	Gas	Dollars
Contract deployment	4697299	$8.63	Accuse$(\cdot)$	223766	$0.41
Deposit$(\cdot)$	105436	$0.19	CheckCircuits$(\cdot)$	66991+	$0.12+
PostRequests$(\cdot)$	405657	$0.74	VerifyExchange$(\cdot)$	61822	$0.11
SubmitResponse$(\cdot)$	97400	$0.18	VerifyGeneralFunc$(\cdot)$	275279	$0.51
ClaimServiceFee$(\cdot)$	33103	$0.06	zkVerify$(\cdot)$	2286423	$4.20

cally expensive since it essentially requires compiling smart contracts on-chain. Therefore, we treat the function circuit provided by the accuser as non-trivial by default and only verify its non-triviality if the accused indicates its triviality before providing auxiliary information. To verify, **CC** creates an oracle contract and a job verifying the non-triviality of smart contracts on an oracle service, e.g., Chainlink [13]. **CC** creates requests through API calls to the oracle and receives responses. Oracle nodes can perform off-chain control and data flow analysis on EVM (Ethereum Virtual Machine) bytecode of smart contracts via EthIR [3] and Securify [86], among other tools. The costs are initially paid by the complainant and charged to the accuser if the function is indeed trivial.

Zero-knowledge coordinator. We implement a privacy-preserving coordinator at the same source [44] and additionally demonstrate a verifier contract for the special information exchange function with Circom compiler [52]. Plonk [38] is employed for proof generation and verification.

B. Discussion

Limitation of supported functions. Overall we prefer $f(\cdot)$ to be light computations, e.g., the entire or the most significant bit (MSB) of the reconstructed data entry. Because it is uneconomical to verify computation-intensive functions when the privacy worth is not high enough. But we also consider this requirement to be undemanding. No matter what the colluding parties have computed during collusion, at least some party, with arbitrary private knowledge, learns some bit(s) or a fraction of a bit about the queried entry from the output. Otherwise, we consider the collusion to be not successful. If a party learns some bits, it can generate a simple explicit function for verification purposes.

Note that although we focus on $f(\cdot)$ that reveals explicit bits about the index or entry, it is feasible to consider arbitrary information gain, e.g., learning that the index is not 1. We only need to update inequality ① in Proposition 1, which counters false accusations from informed servers. Suppose servers guess the arbitrary information gain correctly with probability $\leq \tilde{p}$. Inequality ① is updated to $(1 - \max\{\tilde{p}, \frac{\tilde{p}\omega+1}{\omega+1}\})\lambda_f > \max\{\tilde{p}, \frac{\tilde{p}\omega+1}{\omega+1}\}\}\lambda_r$, i.e., $\lambda_f > \frac{\tilde{p}\omega+1}{(1-\tilde{p})\omega}\lambda_r$. This means that for larger databases, the small information gain of learning that an index is not a specific one requires higher deposits to discourage false accusations. In such applications, the designer can make a trade-off between the smallest information gain to capture and reasonable λ_f.

Set Δ and Δ^*. If the auxiliary information and proof collection time, Δ and Δ^*, are too small, an accuser can initiate denial-of-service (DoS) attacks on the accused servers. When the accused cannot provide auxiliary information or proof in time, the accuser can force a successful false accusation. To set Δ, we capture the maximum allowed offline time δ_1 and the downtime δ_2 that can be caused by DoS attacks. We then let $\Delta > \max\{\delta_1, \delta_2\}$. To set Δ^*, we let it be larger than δ_1 and δ_2 plus the proof generation time. For proof generation, the least efficient case is multiple accused parties generating proofs with collaborative zkSNARKs. This takes only $\approx 488\mu s$ per constraint for two parties [74].

To decide δ_1, we can follow the convention in blockchain systems and let it be a few weeks. For example, hashlocked time contracts in Lightning Network have a maximum delay of 14 days [72]; Ethereum imposes offline penalties to motivate nodes to stay online, e.g., validators lose ~40% balance for inactivity of ~36 days [35]. To decide δ_2, network-layer DDoS attacks empirically last for short periods, e.g., Cloudflare reports 99% of such attacks last under 3 hours (and 86% under 10 minutes) in the first quarter of 2023 [53]. We can put a conservative estimate of a few weeks for δ_2.

Accommodate small ℓ in infinitely repeated PIR service. To allow $\ell = k$, we can completely replace servers at designated time intervals. Given a finite known service deadline, we achieve the non-collusion equilibrium by Corollary 1.

VII. RELATED WORK

In distributed protocols, participants can collude to compromise various aspects of security, including breaking privacy in multi-party computation (MPC) [41], causing equivocations in state machine replication [79], among many others. This work examines privacy-targeted collusion in PIR, and it is relevant to problems with privacy guarantees including generic MPC.

Mask privacy-targeted collusion. Clarke et al. [22] circumvent collusion problem in private data retrieval through anonymization. As a result, a privacy breach is not bound to an identified entity. It is debatable whether this suffices as an effective means of protecting privacy. Besides, a user needs to function as a node in the network, making it unsuitable for some applications including PIR.

Mitigate privacy-targeted collusion assuming rationality. Mangipudi et al. [65] considers collusion-deterrent information escrow service allowing users to encrypt a message that is only revealed once a pre-specified condition is satisfied. Servers can collude to recover the user message prematurely, but this causes a secret key that can transfer others' deposits to be revealed. There are three major differences from our approach. (1) First, in PIR, servers need to perform computations on queries sent

by users, and encrypting messages in special manners is not compatible with PIR constructions that we consider. (2) Second, we allow servers to have arbitrary prior knowledge and capture information gain smaller than learning the entire query or response, i.e., a nontrivial function $f(\cdot)$ on the user secret. Combined with a finite database, random and false accusations become active concerns. (3) In our setting, individual servers earn the entire worth of user secrets instead of splitting the value V among colluders. Besides, we consider repeated runs of the service and malicious parties.

Deter correctness-targeted collusion. Previous works have explored ensuring correct executions by enforcing accountability [21], [46], [82], where at least some misbehaving parties can be located in case of violations. We refer the readers to pertinent surveys and focus on the game-theoretic approaches here. Two main game-theoretic collusion mitigation techniques have been proposed, *slashing* non-reporting collusion participants [29], [92] and having *undercover police* send fake collusion requests to catch colluders [88].

For the first method, Dong et al. [29] focus on ensuring correctness in replication-based cloud computing and propose a solution involving inducing betrayal among colluding parties. The four major differences from mitigating *privacy-targeted collusion* are that (1) first, the correctness-targeted collusion affects protocol executions, thus explicitly leaving observable traces and making collusion evidence well-defined. (2) Privacy-targeted collusion can happen (repeatedly) long after the protocol has terminated. (3) False accusations from parties with information advantages are not a concern there. But in privacy-targeted collusion, servers may hold arbitrary private knowledge, which needs to be distinguished from learning user secrets via collusion. (4) Privacy protection during collusion deterrence is naturally not required there. Besides, computing parties there collude via a smart contract, and we allow the collusion process to be unobserved.

Yakira et al. [92] presents a slashing mechanism for collusion mitigation in threshold cryptosystems. All participants initially register in an escrow service. One can frame colluding players by submitting proper evidence to the service, which slashes all other agents by burning their deposits and rewarding the reporting agent. There, privacy leakage and active false accusations are not a concern, unlike in PIR where there is a finite number of entries. Besides, the evidence and its verification are not discussed.

Ciampi et al. [20] present collusion preserving secure (CP-secure) computation protocol, utilizing a collateral and compensation mechanism to disincentivize parties from aborting. Parties initially deposit collaterals and can withdraw only if all messages and executions are correct. Otherwise, the collateral is taken from dishonest players to reward others. In this design, correct messaging and executions are also well-defined and self-contained.

In the second approach, Wang et al. [88] consider collusion deterrence in MPC. One game-theoretic approach proposed there and also in [89] is to have undercover agents disguised as corrupted parties to catch colluders. In our context, sending fake collusion requests can reduce servers' confidence in others' compliance, but one needs to assume a trusted party to initiate dummy collusion.

VIII. Concluding remarks and future directions

Following the intuition of a secret-shared bit guessing game, we design and implement a mechanism that discourages participants without network advantages from collusion to recover user secrets in multi-server PIR systems, which repeatedly provide query services to users, assuming a large ℓ. The design itself rejects manipulation from users and servers with private knowledge.

Because for practicality, we desire many servers to collectively provide the PIR service, blockchains and systems alike, where the available number of computing parties is much greater than what is needed for performing computations, become suited application scenarios.

Future directions. Our work opens up an interesting line of research. One intriguing direction for future work is to study mechanisms for other secret-sharing style applications or generic MPC. While the rationale behind our approach applies to many MPC tasks, the analysis demonstrates that many nuanced issues will have to be addressed both in the analysis and implementation of the evidence verification.

References

[1] C. Aguilar-Melchor, J. Barrier, L. Fousse, and M.-O. Killijian, "Xpir: Private information retrieval for everyone," *Proceedings on Privacy Enhancing Technologies*, vol. 2, no. 2016, pp. 155–174, 2016.

[2] I. Ahmad, Y. Yang, D. Agrawal, A. El Abbadi, and T. Gupta, "Addra: Metadata-private voice communication over fully untrusted infrastructure," in *15th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 21)*, 2021.

[3] E. Albert, P. Gordillo, B. Livshits, A. Rubio, and I. Sergey, "Ethir: A framework for high-level analysis of ethereum bytecode," in *ATVA*, 2018, pp. 513–520.

[4] A. Ambainis, "Upper bound on the communication complexity of private information retrieval," in *ICALP*, 1997, pp. 401–407.

[5] S. Angel, H. Chen, K. Laine, and S. Setty, "Pir with compressed queries and amortized query processing," in *2018 IEEE symposium on security and privacy (SP)*. IEEE, 2018, pp. 962–979.

[6] A. Beimel, Y. Ishai, E. Kushilevitz, and J.-F. Raymond, "Breaking the $o(n^{1/(2k-1)})$ barrier for information-theoretic private information retrieval," in *IEEE FOCS*, 2002, pp. 261–270.

[7] A. Beimel, Y. Ishai, and T. Malkin, "Reducing the servers computation in private information retrieval: Pir with preprocessing," in *Annual International Cryptology Conference*, 2000, pp. 55–73.

[8] A. Beimel and Y. Stahl, "Robust information-theoretic private information retrieval," *Journal of Cryptology*, vol. 20, no. 3, pp. 295–321, 2007.

[9] M. Ben-Or, S. Goldwasser, and A. Wigderson, "Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract)," in *ACM STOC*, 1988, pp. 1–10.

[10] Bitnodes, "Bitcoin reachable nodes," https://bitnodes.io, accessed: 2023-08-03.

[11] E. Boyle, N. Gilboa, and Y. Ishai, "Function secret sharing: Improvements and extensions," in *ACM CCS*, 2016, pp. 1292–1303.

[12] R. Branzei, D. Dimitrov, and S. Tijs, *Models in cooperative game theory*. Springer Science & Business Media, 2008, vol. 556.

[13] L. Breidenbach, C. Cachin, B. Chan, A. Coventry, S. Ellis, A. Juels, F. Koushanfar, A. Miller, B. Magauran *et al.*, "Chainlink 2.0: Next steps in the evolution of decentralized oracle networks," 2021.

[14] B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, and G. Maxwell, "Bulletproofs: Short proofs for confidential transactions and more," in *2018 IEEE symposium on security and privacy (SP)*. IEEE, 2018, pp. 315–334.

[15] M. Byali, H. Chaudhari, A. Patra, and A. Suresh, "Flash: Fast and robust framework for privacy-preserving machine learning," *Proceedings on Privacy Enhancing Technologies*, vol. 2, pp. 459–480, 2020.

[16] M. Campanelli, D. Fiore, and A. Querol, "Legosnark: Modular design and composition of succinct zero-knowledge proofs," in *ACM CCS*, 2019, pp. 2075–2092.

[17] H. Chaudhari, R. Rachuri, and A. Suresh, "Trident: Efficient 4pc framework for privacy preserving machine learning," 2022.

[18] Y. M. Chee, T. Feng, S. Ling, H. Wang, and L. F. Zhang, "Query-efficient locally decodable codes of subexponential length," *computational complexity*, vol. 22, no. 1, pp. 159–189, 2013.

[19] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, "Private information retrieval," in *IEEE FOCS*, 1995, pp. 41–50.

[20] M. Ciampi, Y. Lu, and V. Zikas, "Collusion-preserving computation without a mediator," in *2022 IEEE 35th Computer Security Foundations Symposium (CSF)*. IEEE, 2022, pp. 211–226.

[21] P. Civit, S. Gilbert, and V. Gramoli, "Polygraph: Accountable byzantine agreement," in *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2021, pp. 403–413.

[22] I. Clarke, O. Sandberg, B. Wiley, and T. W. Hong, "Freenet: A distributed anonymous information storage and retrieval system," in *Designing privacy enhancing technologies*, 2001.

[23] R. Cleve, "Limits on the security of coin flips when half the processors are faulty," in *ACM STOC*, 1986, pp. 364–369.

[24] Cloudflare, "Ecdsa computation," https://blog.cloudflare.com/ecdsa-the-digital-signature-algorithm-of-a-better-internet/.

[25] H. Corrigan-Gibbs, A. Henzinger, and D. Kogan, "Single-server private information retrieval with sublinear amortized time," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2022, pp. 3–33.

[26] P. Dal Bó and G. R. Fréchette, "The evolution of cooperation in infinitely repeated games: Experimental evidence," *American Economic Review*, vol. 101, no. 1, pp. 411–429, 2011.

[27] C. Devet, I. Goldberg, and N. Heninger, "Optimally robust private information retrieval," in *21st USENIX Security Symposium (USENIX Security 12)*, 2012, pp. 269–283.

[28] Y. Dodis, A. Elbaz, R. Oliveira, and R. Raz, "Improved randomness extraction from two independent sources," in *RANDOM*, 2004.

[29] C. Dong, Y. Wang, A. Aldweesh, P. McCorry, and A. van Moorsel, "Betrayal, distrust, and rationality: Smart counter-collusion contracts for verifiable cloud computing," in *ACM CCS*, 2017.

[30] Z. Dvir and S. Gopi, "2-server pir with subpolynomial communication," *Journal of the ACM (JACM)*, vol. 63, no. 4, pp. 1–15, 2016.

[31] K. Efremenko, "3-query locally decodable codes of subexponential length," *SIAM Journal on Computing*, vol. 41, no. 6, 2012.

[32] I. Ehrlich and G. S. Becker, "Market insurance, self-insurance, and self-protection," *Journal of political Economy*, vol. 80, no. 4, 1972.

[33] R. C. Ellickson, *Order without law: How neighbors settle disputes*. Harvard University Press, 1991.

[34] Ethereum, "Ethereum," https://github.com/ethereum/yellowpaper.

[35] ——, "Ethereum inactivity penalty," 2023, https://github.com/ethereum/consensus-specs/blob/v1.4.0-beta.4/specs/phase0/beacon-chain.md#rewards-and-penalties.

[36] Etherscan, "Ethereum node tracker," https://etherscan.io/nodetracker, accessed: 2023-08-03.

[37] J. W. Friedman, "A non-cooperative equilibrium for supergames," *The Review of Economic Studies*, vol. 38, no. 1, pp. 1–12, 1971.

[38] A. Gabizon, Z. J. Williamson, and O. Ciobotaru, "Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge," *Cryptology ePrint Archive*, 2019.

[39] N. Gilboa and Y. Ishai, "Distributed point functions and their applications," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2014.

[40] I. Goldberg, "Improving the robustness of private information retrieval," in *IEEE Symposium on S&P*, 2007, pp. 131–148.

[41] O. Goldreich, S. Micali, and A. Wigderson, "How to play any mental game, or a completeness theorem for protocols with honest majority," in *Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali*, 2019, pp. 307–328.

[42] T. Gong, "Feasibility region 3d plot," 2023, https://sites.google.com/view/pirparameterfeasibilityregion/home.

[43] ——, "Graphical demonstration of coalition size bound," 2023, https://www.desmos.com/calculator/qkphmd9gj2.

[44] ——, "Solidity contracts," 2023, https://github.com/haas256/PIR.

[45] J. Groth, "On the size of pairing-based non-interactive arguments," in *Annual international conference on the theory and applications of cryptographic techniques*, 2016, pp. 305–326.

[46] A. Haeberlen and P. Kuznetsov, "The fault detection problem," in *International Conference On Principles Of Distributed Systems*, 2009, pp. 99–114.

[47] S. M. Hafiz and R. Henry, "A bit more than a bit is more than a bit better: Faster (essentially) optimal-rate many-server pir," *PoPETS*, vol. 2019, no. 4, pp. 112–131, 2019.

[48] R. A. Hardy and J. R. Norgaard, "Reputation in the internet black market: an empirical and theoretical analysis of the deep web," *Journal of Institutional Economics*, vol. 12, no. 3, pp. 515–539, 2016.

[49] G. Helmke and S. Levitsky, "Informal institutions and comparative politics: A research agenda," *Perspectives on politics*, vol. 2, no. 4, pp. 725–740, 2004.

[50] R. Henry, A. Herzberg, and A. Kate, "Blockchain access privacy: Challenges and directions," *IEEE Security & Privacy*, vol. 16, no. 4, pp. 38–45, 2018.

[51] A. Henzinger, M. M. Hong, H. Corrigan-Gibbs, S. Meiklejohn, and V. Vaikuntanathan, "One server for the price of two: Simple and fast single-server private information retrieval," in *Usenix Security*, vol. 23, 2023.

[52] iden3, "Circom 2," https://github.com/iden3/circom, 2023.

[53] C. Inc, "Ddos threat report for 2023 q1 by cloudflare," 2023, https://blog.cloudflare.com/ddos-threat-report-2023-q1/.

[54] T. Itoh and Y. Suzuki, "Improved constructions for query-efficient locally decodable codes of subexponential length," *IEICE Transactions on Information and Systems*, vol. 93, no. 2, pp. 263–270, 2010.

[55] D. Johnson, A. Menezes, and S. Vanstone, "The elliptic curve digital signature algorithm (ecdsa)," *IJIS*, vol. 1, no. 1, pp. 36–63, 2001.

[56] C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan, "GAZELLE: A low latency framework for secure neural network inference," in *27th USENIX Security Symposium*, 2018, pp. 1651–1669.

[57] B. Knott, S. Venkataraman, A. Hannun, S. Sengupta, M. Ibrahim, and L. van der Maaten, "Crypten: Secure multi-party computation meets machine learning," *Advances in Neural Information Processing Systems*, vol. 34, pp. 4961–4973, 2021.

[58] D. M. Kreps and R. Wilson, "Sequential equilibria," *Econometrica: Journal of the Econometric Society*, pp. 863–894, 1982.

[59] S. Kumar, E. Rosnes, and A. G. i Amat, "Private information retrieval in distributed storage systems using an arbitrary linear code," in *ISIT*, 2017, pp. 1421–1425.

[60] E. Kushilevitz and R. Ostrovsky, "Replication is not needed: Single database, computationally-private information retrieval," in *IEEE FOCS*, 1997, pp. 364–373.

[61] L. Lamport, R. E. Shostak, and M. C. Pease, "The byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, 1982.

[62] D. V. Le, L. T. Hurtado, A. Ahmad, M. Minaei, B. Lee, and A. Kate, "A tale of two trees: One writes, and other reads," *PoPETS*, vol. 2020, no. 2, 2020.

[63] W.-K. Lin, E. Mook, and D. Wichs, "Doubly efficient private information retrieval and fully homomorphic ram computation from ring lwe," in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, 2023, pp. 595–608.

[64] G. J. Mailath and L. Samuelson, *Repeated games and reputations: long-run relationships*. Oxford university press, 2006.

[65] E. V. Mangipudi, D. Lu, A. Psomas, and A. Kate, "Collusion-deterrent threshold information escrow," in *2023 IEEE 36th Computer Security Foundations Symposium (CSF)*. IEEE, 2023, pp. 584–599.

[66] S. Matetic, K. Wüst, M. Schneider, K. Kostiainen, G. Karame, and S. Capkun, "BITE: Bitcoin lightweight client privacy using trusted execution," in *USENIX Security*, 2019, pp. 783–800.

[67] S. J. Menon and D. J. Wu, "Spiral: Fast, high-rate single-server pir via fhe composition," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 930–947.

[68] P. Mittal, F. G. Olumofin, C. Troncoso, N. Borisov, and I. Goldberg, "Pir-tor: Scalable anonymous communication using private information retrieval," in *20th USENIX Security Symposium*, 2011.

[69] P. Mohassel and Y. Zhang, "Secureml: A system for scalable privacy-preserving machine learning," in *2017 IEEE Symposium on Security and Privacy (SP)*, 2017, pp. 19–38.

[70] M. H. Mughees, H. Chen, and L. Ren, "Onionpir: Response efficient single-server pir," in *ACM CCS*, 2021, pp. 2292–2306.

[71] R. B. Myerson, *Game theory: analysis of conflict*. Harvard university press, 1997.

[72] L. Network, "Lightning specification," 2019, https://github.com/lightning/bolts/wiki/Lightning-Specification-1.1-Proposal-States#index-of-accepted-proposals.

[73] J. Newsome and D. X. Song, "Dynamic taint analysis for automatic detection, analysis, and signaturegeneration of exploits on commodity software." in *NDSS*, vol. 5. Citeseer, 2005, pp. 3–4.

[74] A. Ozdemir and D. Boneh, "Experimenting with collaborative {zk-SNARKs}:{Zero-Knowledge} proofs for distributed secrets," in *31st USENIX Security Symposium (USENIX Security 22)*, 2022, pp. 4291–4308.

[75] V. Pappas, F. Krell, B. Vo, V. Kolesnikov, T. Malkin, S. G. Choi, W. George, A. Keromytis, and S. Bellovin, "Blind seer: A scalable private dbms," in *IEEE S&P Symposium*, 2014, pp. 359–374.

[76] D. Perito and G. Tsudik, "Secure code update for embedded devices via proofs of secure erasure," in *Computer Security–ESORICS 2010: 15th European Symposium on Research in Computer Security, Athens, Greece, September 20-22, 2010. Proceedings 15*. Springer, 2010, pp. 643–662.

[77] K. Qin, H. Hadass, A. Gervais, and J. Reardon, "Applying private information retrieval to lightweight bitcoin clients," in *Crypto Valley Conference on Blockchain Technology (CVCBT)*, 2019, pp. 60–72.

[78] A. A. Razborov and S. Yekhanin, "An\omega (n^ 1/3) lower bound for bilinear group based private information retrieval," in *IEEE FOCS*, 2006, pp. 739–748.

[79] F. B. Schneider, "Implementing fault-tolerant services using the state machine approach: A tutorial," *ACM Computing Surveys (CSUR)*, vol. 22, no. 4, pp. 299–319, 1990.

[80] R. Selten and R. S. Bielefeld, *Reexamination of the perfectness concept for equilibrium points in extensive games*. Springer, 1988.

[81] N. B. Shah, K. Rashmi, and K. Ramchandran, "One extra bit of download ensures perfectly private information retrieval," in *ISIT*. IEEE, 2014, pp. 856–860.

[82] P. Sheng, G. Wang, K. Nayak, S. Kannan, and P. Viswanath, "Bft protocol forensics," in *Proceedings of the 2021 ACM SIGSAC conference on computer and communications security*, 2021, pp. 1722–1743.

[83] P. Syverson, R. Dingledine, and N. Mathewson, "Tor: The second generation onion router," in *Usenix Security*, 2004, pp. 303–320.

[84] R. Tajeddine, O. W. Gnilke, and S. El Rouayheb, "Private information retrieval from mds coded data in distributed storage systems," *IEEE ToIT*, vol. 64, no. 11, pp. 7081–7093, 2018.

[85] A. W. Tools, "Gwei to usd calculator," https://automatedwebtools.com/usd-eth-gas-fee/, 2023, accessed: 2023-08-03.

[86] P. Tsankov, A. Dan, D. Drachsler-Cohen, A. Gervais, F. Buenzli, and M. Vechev, "Securify: Practical security analysis of smart contracts," in *ACM CCS*, 2018, pp. 67–82.

[87] S. Wagh, D. Gupta, and N. Chandran, "Securenn: 3-party secure computation for neural network training." *Proceedings on Privacy Enhancing Technologies*, vol. 2019, no. 3, pp. 26–49, 2019.

[88] Z. Wang, S.-C. S. Cheung, and Y. Luo, "Information-theoretic secure multi-party computation with collusion deterrence," *IEEE TIFS*, vol. 12, no. 4, pp. 980–995, 2016.

[89] Z. Wang, Y. Luo, and S.-c. Cheung, "Efficient multi-party computation with collusion-deterred secret sharing," in *ICASSP*. IEEE, 2014, pp. 7401–7405.

[90] S. Wehner and R. De Wolf, "Improved lower bounds for locally decodable codes and private information retrieval," in *ICALP*, 2005, pp. 1424–1436.

[91] K. Wüst, S. Matetic, M. Schneider, I. Miers, K. Kostiainen, and S. Čapkun, "Zlite: Lightweight clients for shielded zcash transactions using trusted execution," in *FC*, 2019, pp. 179–198.

[92] D. Yakira, I. Grayevsky, and A. Asayag, "Rational threshold cryptosystems," *arXiv preprint arXiv:1901.01148*, 2019.

[93] A. C. Yao, "Protocols for secure computations," in *FOCS*, 1982, pp. 160–164.

[94] S. Yekhanin, "Towards 3-query locally decodable codes of subexponential length," *Journal of the ACM*, vol. 55, no. 1, pp. 1–16, 2008.

[95] L. F. Zhang and R. Safavi-Naini, "Verifiable multi-server private information retrieval," in *International Conference on Applied Cryptography and Network Security*. Springer, 2014, pp. 62–79.

APPENDIX A
BLIND COLLUSION PROBLEM

Forging information in collusion. When colluding parties do not verify the data shared by colluding partners, one interesting scenario is where a server prefers to retrieve the index alone or desires to hide its input to avoid being reported. We are interested in the *probability of a colluding server fabricating information and fooling others*, i.e., other colluders mistake the provided fake data as correct, which we denote as $p^\dagger$ hereafter.

A cheating player guesses an index and tries to redirect the collusion protocol to output a targeted index. If the underlying PIR protocol ensures that fabricating information (a query or response) in collusion is computationally hard, then p_f, the probability of creating a valid fake input, is 0. This would mean $p^\dagger = 0$, and we are back to the discussion we have in the main body. We consider the worst case $p_f = 1$. This means that as long as a server can picture a redirecting strategy, it can always fabricate such an input.

Analysis. We first discuss $k = 2$. Let server S_i receive query q_1 and S_j ($i \neq j$) receive query q_2 such that $\mathbb{D}_a$ is the queried entry. To form a redirecting strategy, a server first guesses the intended query. One can guess $\mathbb{D}_a$ correctly with probability $\mathbf{x}_a$ where $\mathbf{x}_a$ is the probability that $\mathbb{D}_a$ is queried. Then the server creates a fake input to redirect the output to another index. We build the intuition with the naive xor-based PIR. We let $q_1 = (j_1, j_2, \ldots, j_m)$, $q_2 = (j_1, j_2, \ldots, j_m, a)$, and summarize the strategies in Table II.

TABLE II: Redirecting strategy summary.

Strategy	S_i	S_j
Amicable	$j_1, j_2, \ldots, j_m$	$j_1, j_2, \ldots, j_m, a$
$j_m \to a$	$j_1, j_2, \ldots, j_{m-1}, a$	$j_1, j_2, \ldots, j_{m-1}$
$j_m \to j_{m+1}$	$j_1, j_2, \ldots, j_{m-1}, j_{m+1}$	$j_1, j_2, \ldots, j_{m-1}, a, j_{m+1}$
$j_m \to j_1$	$j_2, \ldots, j_{m-1}$	$j_2, \ldots, j_{m-1}, a$
$a \to j_m$	$j_1, j_2, \ldots, j_{m-1}, a$	$j_1, j_2, \ldots, j_{m-1}$
$a \to j_{m+1}$	$j_1, j_2, \ldots, j_m, a, j_{m+1}$	$j_1, j_2, \ldots, j_m, j_{m+1}$
$j_{m+1} \to j_{m+2}$	$j_1, \ldots, j_m, j_{m+1}, j_{m+2}$	$j_1, \ldots, j_m, a, j_{m+1}, j_{m+2}$
$j_{m+1} \to j_m$	$j_1, \ldots, j_m, j_{m+1}$	$j_1, \ldots, j_{m-1}, a, j_{m+1}$
$j_{m+1} \to a$	$j_1, \ldots, j_m, a, j_{m+1}$	$j_1, \ldots, j_m, j_{m+1}$

Intuitively, if the outputs after the collusion protocol contain more than one index or gibberish when a server itself provides correct inputs, this server knows that the colluding party is cheating. Otherwise, if the collusion output contains only one index, the server may believe that collusion is successful. To find out this probability, we summarize four scenarios where a single index is outputted in Table III.

Strategies	Possible single-index output	Probability $p^\dagger$
(A,A)	a	1
(A,D)	$\{j_m, j_{m+1}\}$	$\mathbf{x}_a(2-\mathbf{x}_a)$
(D,A)	$\{j_m, j_{m+1}\}$	$\mathbf{x}_a(2-\mathbf{x}_a)$
(D,D)	$\{j_m, j_{m+1}, j_{m+2}, a, j_1\}$	$\sum_i \sum_{j\neq i} 4\mathbf{x}_i^2\mathbf{x}_j^2 + \sum_{i\neq a}\sum_{j\neq i,a} 4\mathbf{x}_a\mathbf{x}_i^2\mathbf{x}_j$

TABLE III: Two-server collusion with single-index output.

#D	Probability $p^\dagger$
0	1
1	$\mathbf{x}_a(2-\mathbf{x}_a)$
2	$\sum_i \sum_{j\neq i} 4\mathbf{x}_i^2\mathbf{x}_j^2 + \sum_{i\neq a}\sum_{j\neq i,a} 4\mathbf{x}_a\mathbf{x}_i^2\mathbf{x}_j$
$2 < k' \leq k$	$p_{k'}^\dagger$

TABLE IV: k-server collusion with single-index output.

Row 1 is straightforward: when both servers input the correct information, index a is retrieved with probability 1. In one cheating player case in rows 2 and 3, a single index is outputted when the cheating party either guesses the index correctly (with probability $\mathbf{x}_a$) or guesses the index incorrectly (with probability $(1 - \mathbf{x}_a)$) but redirects the result to $\mathbb{D}_a$ (with probability $\mathbf{x}_a$). And $\mathbf{x}_a + (1 - \mathbf{x}_a)\mathbf{x}_a = \mathbf{x}_a(2 - \mathbf{x}_a)$, which we will denote as $p_1^\dagger$. Index a is never output because the only cheating player either flips a or produces more than one index in the output. The cheating player can always recover a through deduction.

In the case of two cheating players, by observing Table II, we notice that when they have the same strategy or opposite strategies, they recover a. Because they cancel out each other's actions. This happens with probability $p_{2,1}^\dagger = \sum_i \sum_{j\neq i} 4\mathbf{x}_i^2\mathbf{x}_j^2$. For other cases of receiving a single output, $j_y \neq a$ becomes the output if one applies strategy $j_x \to a$ or $a \to j_x$ and the other cheating player plays a strategy $j_y \to j_x$ or $j_x \to j_y$. This occurs with probability $p_{2,2}^\dagger = \sum_{i\neq a} \sum_{j\neq i,a} 4\mathbf{x}_a\mathbf{x}_i^2\mathbf{x}_j$. To sum up, index a is returned with probability $p_{2,1}^\dagger$ while a single index is outputted with probability $p_2^\dagger = p_{2,1}^\dagger + p_{2,2}^\dagger$.

Now we can extend our discussion to a more general $k > 2$ case. When all servers play amicable, a is output with probability 1. When one server cheats, the collusion protocol outputs a single index with probability $\mathbf{x}_a(2-\mathbf{x}_a)$. Similarly, index a is never directly returned, and the only cheating player can always make successful accusations through deductions. When two servers cheat, we have the same probability calculation as before. When $2 < k' \leq k$ servers cheat, they have a single output when each index appears in an even number of strategies, i.e., they have the same or opposite strategies, which occurs with probability $\sum_{i_1} \sum_{i_2 \neq i_1} \cdots \sum_{i_{k'} \neq i_{k'-1}} 2^{k'}\mathbf{x}_{i_1}^2\mathbf{x}_{i_2}^2\cdots\mathbf{x}_{i_{k'}}^2$ or when a and exactly one other index $b \neq a$ appear in an odd number of strategies while every other index appears in an even number of strategies, which we compute next. Let n_a, n_b be the number of appearances of a and b in strategies. We know that $n_a = 1, 3, \ldots, 2\lceil\frac{k'}{2}\rceil - 1$, $n_b = 1, 3, \ldots, 2\lceil\frac{2k'-n_a}{2}\rceil - 1$ and $n_b \leq k'$. Let $n_c = 2k' - n_a - n_b$. For each k', we have

$$p_{k'}^\dagger = \sum_{i_1} \sum_{i_2 \neq i_1} \cdots \sum_{i_{k'} \neq i_{k'-1}} 2^{k'}\mathbf{x}_{i_1}^2\mathbf{x}_{i_2}^2\cdots\mathbf{x}_{i_{k'}}^2 + \sum_{n_a} 2^{k'}\mathbf{x}_a^{n_a} \cdot$$

$$\sum_{n_b}\mathbf{x}_b^{n_b} \sum_{i_1 \neq a,b}\sum_{i_2 \neq i_1,a,b} \cdots \sum_{i_{n_c} \neq i_{n_c-1},a,b} \mathbf{x}_{i_1}^2\mathbf{x}_{i_2}^2\cdots\mathbf{x}_{i_{n_c}}^2$$

In these scenarios with two or more cheating players, cheating servers cannot distinguish between different scenarios with a single output.

A. Implications for the sequential game

Now G is a game of incomplete information: in round three, an agent is not certain of which actions the other agent has taken. Backward induction no longer applies, and we solve for sequential equilibrium instead.

Information sets. An *information set* for a player S_i contains all decision nodes, at which player S_i takes an action, indistinguishable by S_i. Many information sets in our game are trivial, i.e., containing a single node. For example, in round 1, both agents are aware of the other agent's interests in collusion. However, there are non-trivial information sets, e.g., the ones prefixed "12" in Figure 2, where the two make simultaneous decisions.

Solution concept. We consider another refinement of Nash equilibrium for sequential games with incomplete information, namely *Sequential Equilibrium* [58]. A sequential equilibrium not only prescribes a strategy for each player but also a *belief* about other players' strategies. To formalize the notion of sequential rationality, we first formalize the notion of a *belief*.

Definition 2. *For each information set I, a belief assessment b gives the conditional probability distribution $b(\cdot|I)$ over all decision nodes $v \in I$.*

For example, $b(D_2|I_1^A)$ is the belief of player S_1 that S_2 takes action D in round one, given that S_1 is in information set $I_1^A = \{(A,A), (A,D)\}$. A belief assessment expresses all such beliefs for all players at their information sets.

Definition 3. *A strategy profile and belief assessment pair $(\mathbf{s}, b)$ is a sequential equilibrium if $(\mathbf{s}, b)$ is sequentially rational and b is consistent with $\mathbf{s}$.*

Given $(\mathbf{s}, b)$, $\mathbf{s}$ is called *sequentially rational* if at each information set I, the player to take an action maximizes her expected payoff, given beliefs $b(\cdot|I)$ and that other players follow $\mathbf{s}$ in the continuation (remaining) game. Given a strategy profile $\mathbf{s}$, for any information set I on the path of play of $\mathbf{s}$, $b(\cdot|I)$ are *consistent* with $\mathbf{s}$ if and only if they are derived using Bayes' rule. For example, let player 1 be at an information set $I = \{x, y\}$. Suppose according to $\mathbf{s}$, x, y are respectively reached with probability $0.2, 0.3$. Then a consistent belief for player 1 is that $b(x|I) = \frac{0.2}{0.2+0.3} = 0.4$ and $b(y|I) = 0.6$.

We address the information sets in rounds one and two of G as information set 1, 2. They are singletons. In round three, we call the information set where one thinks others have played in round two as 3, the information set where an amicable player recognizes that others have played D as 4, and the information set where a deceitful player recognizes that others have played D as 5. We denote reporting (R) at information sets $3, 4, 5$

as R, R_4, R_5. We have the following main theorem for the equilibrium of the collusion game without input verification. In a single-shot game, same as before, we let $\lambda_r = 0$. By enforcing $\lambda_s > 0, \lambda_f > 0$, we discourage false accusations.

Theorem A.1. *In* 2-*party collusion game* G *without input verification in* $(\ell, 2, 1)$-*PIR with* $\theta\ell$ *malicious servers, with probability* $1 - 2^{-\eta}$, *the pair* $(\mathbf{s}, b)$ *of strategy profile* $\mathbf{s} = (\bar{C}DR\bar{R}_4R_5, \ldots)$ *for all servers and belief assessment* b *is the sequential equilibrium if* $\lambda_p > 0$ *where* b *satisfies* $b(D_{\mathbf{s}_{-i}}|I_3^A) = b(D_{\mathbf{s}_{-i}}|I_3^D) = b(R_{\mathbf{s}_{-i}}|I_3^A) = b(R_{\mathbf{s}_{-i}}|I_3^D) = 1$, *and* θ *satisfies* $S^\star$.

Here $b(D_{\mathbf{s}_{-i}}|I_3^A)$ is the belief that others have played D in round two after one arrives at information set 3 by playing A, and $b(R_{\mathbf{s}_{-i}}|I_3^A)$ is the belief that others play R in round three when they all choose an action.

Proof. With probability $1 - 2^{-\eta}$, at least two queried servers are rational. Let the two queried servers be S_1, S_2. We depict the game tree in Figure 6. At information set 4 and 5, the dominant strategies are respectively $\bar{R}_4$ and R_5. We then examine information set 3. It consists of four nodes. From top to bottom, the four nodes are reached with probability $\alpha_2\beta_2, \alpha_2(1-\beta_2)p_1^\dagger, (1-\alpha_2)\beta_2p_1^\dagger, (1-\alpha_2)(1-\beta_2)p_2^\dagger$. Without loss of generality, we examine from one queried party S_1's perspective. By Bayes' formula, when S_1 plays A at information set 2, it believes that it arrives at the information set 3 along move (A, A) or A_1A_2 with probability

$$b(A_2|I_3^A) = \frac{\beta_2}{\beta_2 + (1-\beta_2)p_1^\dagger}$$

Similarly, it believes it arrives at information set 3 along move A_1D_2 with probability $b(D_2|I_3^A) = \frac{(1-\beta_2)p_1^\dagger}{\beta_2+(1-\beta_2)p_1^\dagger}$. When S_1 plays D at information set 2, we can calculate $b(A_2|I_3^D) = \frac{\beta_2p_1^\dagger}{\beta_2p_1^\dagger+(1-\beta_2)p_2^\dagger}$ and $b(D_2|I_3^D) = \frac{(1-\beta_2)p_2^\dagger}{\beta_2p_1^\dagger+(1-\beta_2)p_2^\dagger}$. For visual clarity, we denote the four probabilities as x_1, x_2, y_1, y_2. If S_1 arrives at information set 3 after playing D, the expected returns are:

$$y_1\Big(\alpha_4\beta_3(V + \lambda_s + \lambda_r) + \alpha_4(1-\beta_3)(V + \lambda_s + \lambda_r) +$$
$$(1-\alpha_4)\beta_3(V + \lambda_s) + (1-\alpha_4)(1-\beta_3)(V + \lambda_s)\Big) +$$
$$y_2\Big(\alpha_4\beta_4(\lambda_s + p_{2,1}^\dagger(V - \frac{\lambda_p - \lambda_r}{2} - \lambda_s)) + \alpha_4(1-\beta_4)\cdot$$
$$(\lambda_s + p_{2,1}^\dagger(V + \lambda_r)) + (1-\alpha_4)\beta_4(\lambda_s + p_{2,1}^\dagger(V - \lambda_p - \lambda_s)) +$$
$$(1-\alpha_4)(1-\beta_4)(\lambda_s + p_{2,1}^\dagger V)\Big) = \alpha_4\big(y_1\lambda_r +$$
$$y_2 p_{2,1}^\dagger(\beta_4 \frac{\lambda_p + \lambda_r}{2} - \lambda_r)\big) + \epsilon = \alpha_4 y_2 p_{2,1}^\dagger \beta_4 \frac{\lambda_p}{2} + \epsilon$$

Here ϵ is a term that contains only βs and does not depend on S_1's strategies. Since $\lambda_p > 0$, $\alpha_4 = 1$ maximizes the expected return. Similarly, $\beta_4 = 1$. Now we solve for α_3 and β_3. Consider S_1 arriving at the information set 3 after playing A. There are three possible supports to consider for S_1: $\{R\}, \{\bar{R}\}, \{R, \bar{R}\}$ (randomization over $R, \bar{R}$). To make R sequentially rational, we need the expected returns from playing R to be higher than playing $\bar{R}$:

$$R: \ x_1\big(\beta_3(V - \frac{\lambda_p - \lambda_r}{2}) + (1-\beta_3)(V + \lambda_s + \lambda_r)\big) - x_2\lambda_p$$
$$\bar{R}: \ x_1\big(\beta_3(V - \lambda_p) + (1-\beta_3)(V + \lambda_s)\big) - x_2\lambda_p$$

Then playing R generates higher returns than $\bar{R}$, and $\alpha_3 = 1$. Similarly, $\beta_3 = 1$. We next solve for α_2 and β_2. At information set 2, there are three possible supports to consider for S_1: $\{A\}, \{D\}, \{A, D\}$. To make D sequentially rational, then we need the expected returns from playing D to be higher than playing A. The simplified expected returns from playing the two moves are as follows:

$$A: \beta_2(V - \frac{\lambda_p - \lambda_r}{2}) + (1-\beta_2)\big(p_1^\dagger(-\lambda_p) + (1-p_1^\dagger)(-\lambda_p)\big)$$
$$= \beta_2(V - \frac{\lambda_p}{2}) - (1-\beta_2)\lambda_p$$
$$D: \beta_2(V + \lambda_s + \lambda_r) + (1-\beta_2)\big(p_2^\dagger p_{2,1}^\dagger(V - \frac{\lambda_p - \lambda_r}{2} - \lambda_s)$$
$$+ \lambda_s\big)$$

The first quantity is smaller than the second quantity since both terms are smaller. Because $\lambda_p > 0$ and $p_2^\dagger p_{2,1}^\dagger V + (1 - \frac{p_2^\dagger p_{2,1}^\dagger}{2})\lambda_p + (1 - p_2^\dagger p_{2,1}^\dagger)\lambda_s > 0$. This means that $\alpha_2 = 0$. Likewise, $\beta_2 = 0$. Then $x_1 = y_1 = 0, x_2 = y_2 = 1$, and the above is consistent with the beliefs. In round one, playing $\bar{C}$ generates the same returns as CD. Then the sequential equilibrium is all players playing $\bar{C}DR\bar{R}_4R_5$. $\square$

APPENDIX B
MULTI-SERVER PIR CONSTRUCTIONS

A. 2-*server PIR*

We first describe two related primitives, Function Secret Sharing (FSS) and its special case Distributed Point Function (DPF). An m-party Function Secret Sharing (FSS) scheme for a function family $\mathcal{F}$ from $\{0,1\}^n$ to an Abelian group $\mathbb{G}$ comprises key generation function $Gen(\cdot)$ and an evaluation function $Eval(\cdot)$. For a function $f \in \mathcal{F}$, algorithm $Gen(\cdot)$ returns m-tuple keys $k_1, ..., k_m$ where each key k_i defines function $f_i(x) = Eval(i, k_i, x)$ so that $f(x) = \sum_{i=1}^m f_i(x)$. With less than m keys, it is computationally hard to derive f. Distributed Point Function (DPF) [39] is FSS for point functions. A point function $f_{a,b} : \{0,1\}^n \to G$ for $a \in \{0,1\}^n, b \in \mathbb{G}$ evaluates to b on input a and 0 otherwise. **Boyle-Gilboa-Ishai's construction for** 2-**server cPIR.** This construction utilizes a PRG-based 2-party DPF. The client distributes the point function $f_{a,1} : [N] \to \mathbb{Z}_2$ to two servers to retrieve $\mathbb{D}_a$. She generates a pair of keys (k_1, k_2) with a secure PRG, obtains function shares (f_1, f_2) and sends (f_i, k_i) to server i. Server i responds with $\sum_{j=1}^N \mathbb{D}_j f_i(j)$. The client then obtains $\mathbb{D}_a$ by xor-ing two responses.

B. *k-server PIR*

We now work in an arbitrary finite field $\mathbb{F}$ and represent each entry in $\mathbb{F}$. Suppose it takes e elements to describe each entry. Then $\mathbb{D}$ can be represented in $\mathbb{F}$ as an $e \times N$ matrix.
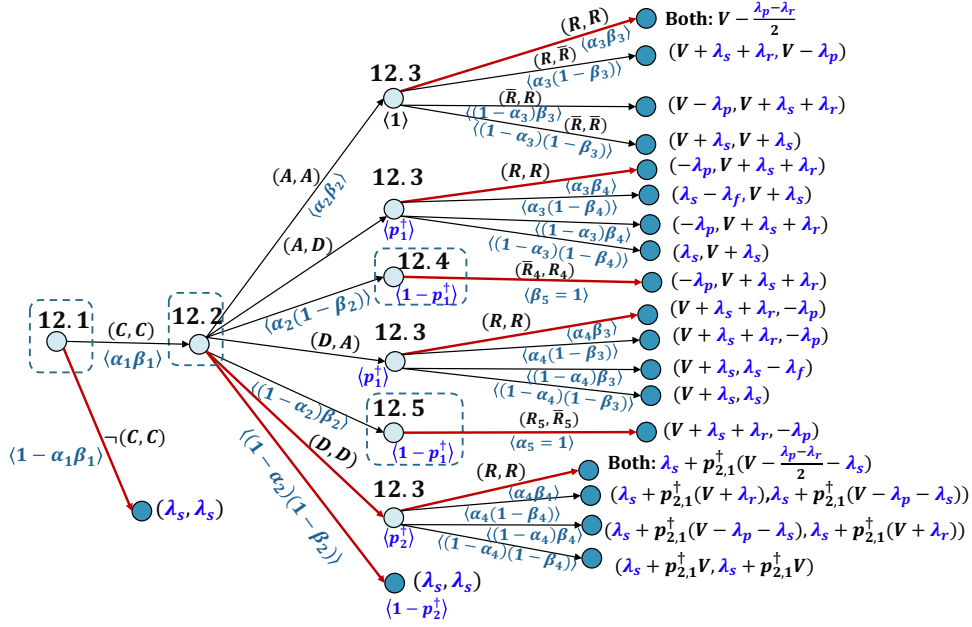
Fig. 6: Extensive form of G without verification ($k = 2$).

We denote each entry as a vector $\mathbb{D_i}$ in this section instead of $\mathbb{D}_i$ to signal this e-element representation. $\mathbb{D_i}_j$ locates the j-th element of $\mathbb{D_i}$.

Hafiz-Henry's construction for k-server cPIR. We present the protocol for querying $k = 2^K$ ($K \in \mathbb{Z}^+$) servers. The client intends to fetch $\mathbb{D}_a$ and needs to distribute the point function $f_{a,1} : [N] \to \mathbb{Z}_2$ to k servers. She generates K independent (2,2)-DPF key pairs, $(\mathsf{k}_0^{(0)}, \mathsf{k}_1^{(0)}), ..., (\mathsf{k}_0^{(K-1)}, \mathsf{k}_1^{(K-1)})$, for the point function $f_{a,1}$. The client now sends to each server j the key share $(\mathsf{k}_{j_{K-1}}^{(K-1)}, ..., \mathsf{k}_{j_0}^{(0)})$. Each server j expands each of these keys into a length-N vector of bits and then concatenates the K vectors component-wise to obtain a length-N vector $\mathbf{v}$ of K-bit strings. Server j then goes through $\mathbf{v}$ component by component, xor-ing the $\mathbf{v}[i]$th word of $\mathbb{D}_i$ into a running total when $\mathbf{v}[i] = 1$. By construction, the a-th bit produced by each DPF key pair differs and all others are equal; thus, the XOR that each server produces is identical up to but not including which word of $\mathbb{D}_a$ it includes; moreover, one server includes no word of $\mathbb{D}_a$ at all. Taking this latter server's response and xor-ing it with the responses from each of the other servers yields each of the words comprising $\mathbb{D}_a$ (in some random, but known to the querier, order).

For performance evaluation in [47], the experiments are conducted on a workstation running Ubuntu 18.04.2 LTS on an AMD Ryzen 7 2700x eight-core CPU @ 4.30 GHz with 16 GiB of RAM and a 1 TB NVMe M.2 SSD. According to the results therein, sampling DPF keys on the clients' side takes less than 100 microseconds, and key expansion on the server side takes about 0.5ms for a database with 2^{20} rows. The answer reconstruction time is about constant for each extra bit retrieved. To reconstruct an entry of 1 GiB, the reconstruction time is about 130 ms.

Efficiency comparison. SealPIR [5] has $O(\sqrt{N})$ communica-tion complexity, and incurs $\approx$ 100x response size overhead and heavy computation (e.g., 400s for a 30 KB entry with $N = 2^{20}$); OnionPIR [70] has $O(\log N)$ communication costs and results in 4.2x response size overhead and similar computation costs; The Corrigan-Gibbss et al. proposal [25] has $O(\sqrt{N})$ com-munication costs and achieves amortized sublinear server-side computation time by having the client download a one-time hint of size sublinear in $|\mathbb{D}|$; SimplePIR [51] achieves 6.5 GiB/s/core throughput but requires $O(\sqrt{N})$ communication costs, e.g., 124 MB hint plus 242 KB per query communication for a 1 GiB database. DEPIR [63] has poly $log(n)$ communication complexity but incurs server storage superlinear in the database size and requires ideal obfuscation.

For TEE-based solutions, the processing time for Bite [66] (oblivious database) is around 0.5s for incoming client requests and around 79s for new blocks. T^3 [62] consumes about 0.64-2.43ms for smaller Oblivious RAM block sizes and read-only accesses. The two multi-server PIR schemes are more efficient.

APPENDIX C
SELF-INSURANCE

A single server S_e is exiting. Suppose users arrive at rate $\frac{\Omega}{T}$: during each time unit, there are $k\Omega$ queries. S_e expects to receive $\frac{k\Omega}{\ell}$ queries and can recover all $\frac{k\Omega}{\ell}$ secrets in the worst case. The amount of accumulated service fees in those instances are $\sum_{t=1}^{T} k \frac{k\Omega}{\ell} \lambda_s (1 + r)^{T-t}$ in expectation. The expected amount of fines needed is $\sum_{t=1}^{T} \frac{k-1}{k} \frac{k\Omega}{\ell} \lambda_p (1-r')^{T-t}$. We continue to let the ratio between inexecutable fines and service fees be ≤ 1:

$$\frac{\sum_{t=0}^{T} \frac{k-1}{k} \frac{k\Omega}{\ell} \lambda_p (1-r')^{T-t} - \lambda_p}{\sum_{t=1}^{T} \frac{k\Omega}{\ell} \lambda_s (1+r)^{T-t}} \leq 1$$

Multiple servers are exiting. Suppose $m \geq 2$ servers are leaving the service, and users arrive at rate $\frac{\Omega}{T}$. Let $b = \min\{m, k\}$, i servers ($i = b, b-1, \ldots, 1$) are queried at the same time with probability $\frac{\binom{m}{i}\binom{\ell-m}{k-i}}{\binom{\ell}{k}}$. Expected inexecutable fines amount to $\sum_{t=0}^{T} \sum_{i=1}^{b} \frac{\binom{m}{i}\binom{\ell-m}{k-i}}{\binom{\ell}{k}} \frac{k-i}{k} \Omega i \lambda_p (1-r')^{T-t} - m\lambda_p$ and expected accumulated service fees are $\sum_{t=0}^{T} \sum_{i=1}^{b} \frac{\binom{m}{i}\binom{\ell-m}{k-i}}{\binom{\ell}{k}} \Omega k \lambda_s (1+r)^{T-t}$. We then need

$$\frac{\sum_{t=0}^{T} \sum_{i=1}^{b} \frac{\binom{m}{i}\binom{\ell-m}{k-i}}{\binom{\ell}{k}} \frac{k-i}{k} \Omega i \lambda_p (1-r')^{T-t} - m\lambda_p}{\sum_{t=0}^{T} \sum_{i=1}^{b} \frac{\binom{m}{i}\binom{\ell-m}{k-i}}{\binom{\ell}{k}} \Omega k \lambda_s (1+r)^{T-t}} \leq 1$$