Cryptographic Randomized Response Techniques

Andris Ambainis[1], Markus Jakobsson[2], and Helger Lipmaa[3]

[1] CS Division, University of California, Berkeley, CA 94720, USA
ambainis@cs.berkeley.edu
[2] RSA Laboratories, 174 Middlesex Turnpike, Bedford, MA 01730, USA
mjakobsson@rsasecurity.com
[3] Laboratory for Theoretical CS, Department of CS&E
Helsinki University of Technology, P.O.Box 5400, FIN-02015 HUT, Espoo, Finland
helger@tcs.hut.fi

Abstract. We develop cryptographically secure techniques to guarantee unconditional privacy for respondents to polls. Our constructions are efficient and practical, and are shown not to allow cheating respondents to affect the "tally" by more than their own vote—which will be given the exact same weight as that of other respondents. We demonstrate solutions to this problem based on both traditional cryptographic techniques and quantum cryptography.
Keywords: binary symmetric channel, oblivious transfer, polling, privacy, privacy-preserving data-mining, randomized response technique

1 Introduction

In some instances, privacy is a matter of keeping purchase information away from telemarketers, competitors, or other intruders. In other instances, privacy translates to security against traffic analysis, such as for web browsing; or to security of personal location information. In still other instances, which we study in this paper, privacy is a *precondition* to being able to obtain answers to important questions. Two concrete examples of instances of latter are *elections* and *surveys/polls*.

While the first of these examples is the one of the two that has received—by far—the most attention in the field of cryptography, there are important reasons to develop better privacy tools for polling. Surprisingly, the two examples (namely, elections and polls), while quite similar at a first sight, are very different in their requirements. Since it is typically the case that there is more funding available for providing privacy in elections than in surveys and polls, it follows that the tallying process in the former may involve more costly steps than that in the latter—whether the process is electronic (using, e.g., mix networks) or mechanic. Second, while in the case of the voting scheme, we have that users need to entrust their privacy with some set of authorities, it is often the case that there is less trust established between the parties in polls. Yet another reason to treat the two situations separately is that elections involve many more respondents than polls typically do, thereby allowing a unique opinion (e.g., vote) to be hidden among many more in the case of elections than in the case of polls. Finally, while elections require as exact tallying as is possible, *statistical truths* are both sufficient and desirable in polls. This allows the use of polling techniques that are very different from election techniques—in terms of their cost; how tallying is done; and how privacy is protected.

F. Bao et al. (Eds.): PKC 2004, LNCS 2947, pp. 425–438, 2004.

While not given much attention in cryptography, important work on polling has been done in statistics. In particular, the *randomized response technique* (RRT) was proposed by Warner [War65] in 1965, with the goal of being used in polls relating to sensitive issues, such as drug abuse, sexual preferences and shoplifting. The underlying idea behind Warner's proposal (alternative RRTs have been proposed since then) is for respondents to randomize each response according to a certain, and known, probability distribution. More precisely, they answer the question truthfully with some probability $p_{\mathsf{ct}} > 1/2$, while with a fixed and known probability $1 - p_{\mathsf{ct}}$ they lie. Thus, users can always claim that their answer—if it is of the "incriminating" type—was a lie. When evaluating all the answers of the poll, these lies become statistically insignificant given a large enough sample (where the size of the sample can be simply computed from the probability distribution governing lying.)

However, a pure RRT by itself is not well suited for all types of polls. E.g., it is believed that people are more likely to vote for somebody who leads the polls than somebody who is behind. Therefore, it could be politically valuable not to lie (as required by the protocol) in polls relating to one's political opinion, and therefore have one's "vote" assigned a greater weight. (This is the case since people with the opposite opinion—if honestly following the protocol—will sometimes cast a vote according to your opinion, but you would never cast a vote according to their opinion, assuming you are willing to cheat.) While the results of the poll remain meaningful if *everybody* cheats (i.e., tells the truth with a probability different from that specified by the protocol), this is *not* the case when only some people deviate from the desired behavior. Also, while one might say that the increased weight in the polls is gained at the price of the cheater's privacy, this is not necessarily the case if the cheater *claims* to have followed the protocol, and there is no evidence to the contrary.

To address the problem of cheating respondents in RRT, we propose the notion of *cryptographic randomized response technique* (CRRT), which is a modification of RRT that prevents cheating. We present three efficient protocols for CRRT; two of them using classic cryptographic methods (and being efficient for different values of p_{ct}), and one using quantum methods. Importantly, the quantum RRT protocol is implementable by using contemporary technology. We give rigorous proofs of security for one of the classical protocols and for the quantum protocol.

For all of our proposed solutions, the privacy of the respondent will be guaranteed information-theoretically (more precisely, statistically). This is appropriate to stimulate truthful feedback on topics that may affect the respondent for years, if not decades. All proposed solutions also *guarantee* that the respondents reply based on the desired probability distributions. Clearly, this requires that the respondent cannot determine the outcome of the protocol (as viewed by the interviewer) before the end of the protocol. Otherwise, he could simply halt the execution of the protocol to suppress answers in which the communicated opinion was a lie. We will therefore require protocols to offer privacy for the *interviewer* as well as for the respondent, meaning that the respondent cannot learn what the outcome of the protocol is, as seen by the interviewer. (One could relax this requirement slightly to allow the respondent to learn the outcome at the same time as the interviewer does, or afterward.)

While we believe that it is important to prevent the respondent from biasing the outcome by selective halting (corresponding to the protocol being *strongly secure*), we also describe simplified versions of our protocols in which this protection mechanism is not available. Such simplified versions (which we refer to as *weakly secure*) can still be useful in some situations. They may, for example, be used as the default scheme for a given application—where they would be replaced by their strongly secure relatives if too many interactions are halted prematurely. (The decision of when the shift would be performed should be based on standard statistical methods, and will not be covered herein.) The benefit of considering such dual modes is that the weakly secure versions typically are computationally less demanding than the strongly secure versions.

Finally, we also discuss cryptographic enhancements to two alternative RRT techniques. In the first, referred to as RRT-IQ, the respondent always gives the truthful answer to the question he is presented with. However, with a certain probability, he is presented with an Innocous Question instead of the intended question. A second alternative RRT technique is what is referred to as *polychotomous* RRT. In this version of RRT, the respondent is given more than two possible options per question.

Other Applications. Our first protocol uses a novel protocol for information-theoretically secure *verifiable oblivious transfer* that enables easier zero-knowledge proofs on the properties of the transferred values. The new verifiable oblivious transfer protocol may also be useful in other applications. While our main designated application is polling, our techniques have also several other applications, in particular in the privacy-preserving data-mining. They are also related to several fundamental cryptographic problems. For example, our protocols Wagner's technique are also efficient implementations of the *verifiable binary symmetric channel*. (See Section 3.)

New Verifiable Commitment Scheme. One of our RRT protocols uses a novel (and as far as we know, the first) two-round verifiable commitment scheme based on the (non-verifiable) commitment scheme by Naor and Pinkas [NP01]. Verifiable commitment schemes have a huge range of applications.[4]

Outline. We first review the details of the randomized response technique (Section 2), after which we review some related work in cryptography (Section 3). We then introduce the cryptographic building blocks of our protocols (Section 4). We then describe the functionality of our desired solution in terms of functional black boxes and protocol requirements (Section 5). In Section 6, we present our secure CRRT protocols. In Section 7 we describe cryptographic solutions to other variants of the standard RRT. The appendix contains additional information about the new oblivious transfer protocol and about the quantum RRT protocol.

2 Short Review of Randomized Response Technique

When polling on sensitive issues like sexual behavior or tax evasion, respondents often deny their stigmatizing behavior due to the natural concern about their privacy. In

[4] Slightly more efficient and recent verifiable commitment schemes that draw ideas from this paper were proposed by the third author in [Lip03b]. The new schemes can be seamlessly plugged into our first RRT protocol.

1965, Warner [War65] proposed the Randomized Response Technique (RRT) for organization of polls where an unbiased estimator (UE, defined in any standard statistics textbook) to the summatory information —the proportion of people belonging to a stigmatizing group A—can be recovered, while the privacy of every individual respondent is protected statistically. Since then, different variations of the RRT have been proposed in statistics, see [CM88] for a survey. These different variations provide, for example, smaller variance, smaller privacy breaches, optimality under different definitions of privacy, and ability to answer polychotomous questions. Next we will give a short overview of three types of RRT.

RRT-W. In Wagner's original method (RRT-W), the respondents provide a truthful answer to the question "Do you belong to a stigmatizing group A?" with a certain fixed and publicly known probability $p_{\mathsf{ct}} > 1/2$. With probability $1 - p_{\mathsf{ct}}$ they lie—i.e., answer the opposite question. Define π_A to be the true proportion of the population that belongs to A (or whose *type* is $t = 1$). Let p_{yes} be the proportion of "yes" responses in the poll. In RRT-W, the *a priori* probability of getting a "yes" response is $p_{\mathsf{yes}} = p_{\mathsf{ct}} \cdot \pi_A + (1 - p_{\mathsf{ct}})(1 - \pi_A)$. In the case of N players, L of which answer "yes", an UE of p_{yes} is $\widehat{p_{\mathsf{yes}}} = L/N$, the sample proportion of "yes" answers. From this, one can simply compute the unbiased estimator of π_A. This equals $\widehat{\pi_A} = \frac{\widehat{p_{\mathsf{yes}}} - (1-p_{\mathsf{ct}})}{2p_{\mathsf{ct}}-1} = \frac{p_{\mathsf{ct}}-1}{2p_{\mathsf{ct}}-1} + \frac{L}{N} \cdot \frac{1}{(2p_{\mathsf{ct}}-1)}$. Similarly, the variance $\mathrm{var}(\widehat{\pi_A})$ and its UE can be computed.

RRT-IQ. An alternative RRT is the *innocuous question method* (RRT-IQ), first analyzed in [GASH69]. When using RRT-IQ, the respondent answers the sensitive question with a probability p_{ct}, while with probability $1 - p_{\mathsf{ct}}$ to an unrelated and innocuous question, such as "Flip a coin. Did you get tails?". The RRT-IQ achieves the same goals as RRT-W but with less variance [CM88], which makes it more suitable for practical polling. Many other RRT-IQs are known, including some with unknown estimate of the the proportion of the population belonging to the innocuous group.

PRRT. The RRTs for dichotomous polling (where the answer is yes or no) can be generalized to *polychotomous RRT* (PRRT) where the respondent can belong to one of the m mutually exclusive groups $A_1, \ldots, A_m$, some of which are stigmatizing. A typical sensitive question of this kind is "When did you have your first child?", with answers "1—while not married", "2—within 9 months after the wedding" and "3—more than 9 months after the wedding". In many cultures, the answer 1 is stigmatizing, the answer 3 is innocuous, while the answer 2 is somewhere inbetween. The interviewer wants to know an UE for the proportion π_i of people who belong to the group A_i, $i \in [1, m]$. There are many possible PRRTs [CM88, Chapter 3]. One of the simplest is the following technique PRRT-BD by Bourke and Dalenius [CM88]: first fix the probabilities p_{ct} and $p_1, \ldots, p_m$, such that $p_{\mathsf{ct}} + \sum_{i \in [1,m]} p_i = 1$. A respondent either reveals her true type $t \in [1, m]$ with probability p_{ct}, or answers $i \in [1, m]$ with probability p_i. To recover an UE of $\boldsymbol{\pi} := (\pi_1, \ldots, \pi_m)^T$, define $\boldsymbol{p} := (p_1, \ldots, p_m)^T$ and $\boldsymbol{p}_{\mathsf{ans}} = (p_{\mathsf{ans}_1}, \ldots, p_{\mathsf{ans}_m})^T$, where p_{ans_i} is the proportion of people who answer i. Then $\boldsymbol{p}_{\mathsf{ans}} = p_{\mathsf{ct}} \cdot \boldsymbol{\pi} + \boldsymbol{p}$, and hence $\widehat{\boldsymbol{\pi}} = p_{\mathsf{ct}}^{-1} \cdot (\widehat{\boldsymbol{p}_{\mathsf{ans}}} - \boldsymbol{p})$.

3 Related Cryptographic Work

In [KANG99], Kikuchi et al. propose techniques with similar goals as ours. Unaware of the previous work on RRT, the authors reinvent this notion, and propose a protocol for performing the data exchange. However, their protocol is considerably less efficient than ours. Also, it does not offer strong security in our sense. This vulnerability makes their protocol unsuitable for their main application (voting), as well as polls where respondents may wish to bias their answer. Our protocols can be used in their framework.

The cryptographic RRT-W protocol can be seen as an implementation of an *verifiable BSC*, based on either verifiable oblivious transfer or more generally on a suitable commitment scheme. (Protocols for other RRTs implement even more complex channels.) Crépeau and Kilian have showed how to construct (nonverifiable) oblivious transfer protocols and commitment schemes from a (nonverifiable) BSC [CK88, Cré97], but their opposite reductions are less efficient.

There is a very close relationship between our protocols and protocols for oblivious transfer and for the fractional oblivious transfer [BR99]. While our goals are orthogonal to those of oblivious transfer, the techniques are hauntingly similar. In particular, one of our CRRT protocols uses a protocol for oblivious transfer as a building block. While in principle *any* such protocol can be used, it is clear that the properties of the building block will be inherited by the main protocol. Therefore, in order to provide unconditional guarantees of privacy for the respondents, we use a *verifiable* variant of the information theoretic protocol for oblivious transfer, namely that proposed by Naor and Pinkas [NP01]. We leave it as an open question whether the fractional oblivious transfer protocols of [BR99] (that essentially implement verifiable *erasure channel*) can be modified to work in our scenario (where we need to implement verifiable BSC in the case of RRT-W and related information channels without erasure in the case of other RRT protocols) or our protocols can be modified to work in their scenario; at least the first seems clearly not to be the case.

Furthermore, our work is related to the work on Private Information Retrieval (PIR) in that the goal of our interviewer is to retrieve some element from the respondent, without the latter learning what was retrieved. More specifically, if some ℓ out of n elements represent the respondent's opinion, and the remaining $n - \ell$ elements represent the opposite opinion, then the interviewer will learn the respondent's opinion with probability ℓ/n if he retrieves a random element. Of course, in order to guarantee the interviewer that the elements are correctly formed, additional mechanisms are required.

In privacy-preserving data-mining a related data randomization approach has been proposed: namely, the users input their data to the central database (e.g., a loyal customer inputs the name of the product he bought), and the database maintainer needs to do some statistical analysis on the database. However, the maintainer should not be able to recover individual items. Database randomization in the case when the maintainer is limited to the SUM function corresponds exactly to the RRT. For the same reasons as in the RRT, one should not be able to bias the data. Our protocols are also applicable in the privacy-preserving data-mining.

4 Cryptographic Building Blocks

Define $[a,b] := \{a, a+1, \ldots, b-1, b\}$. In the rest of this section we present some cryptographic building blocks that will be used in our CRRT protocols. Throughout this paper, assume that p is a large prime, and q, $q \mid (p-1)$, is another prime. Then $\mathbb{Z}_p^*$ has a unique subgroup G of order q. Let g and h be two generators of G, such that nobody knows their mutual discrete logarithms $\log_g h$ and $\log_h g$. We let k be the security parameter, in our setting we can take $k = q$. In the next two protocols (the Pedersen's commitment scheme and the Naor-Pinkas oblivious transfer protocol), the key K consists of public parameters, $K := (g; h)$.

Pedersen's Commitment Scheme. In this scheme [Ped91], a message $\mu \in \mathbb{Z}_q$ is committed by drawing a random $\rho \leftarrow_R \mathbb{Z}_q$, and setting $\mathsf{C}_K(\mu;\rho) := g^\mu h^\rho$. The commitment can be opened by sending μ and ρ to the verifier. This scheme is *homomorphic*, i.e., $\mathsf{C}_K(\mu;\rho)\mathsf{C}_K(\mu';\rho') = \mathsf{C}_K(\mu+\mu';\rho+\rho')$. Since it is also perfectly hiding and computationally binding, it can be used as a building block in efficient zero-knowledge arguments, such as protocols for arguing the knowledge of plaintext μ.

Verifiable 1-out-of-n Oblivious Transfer. In an $\binom{1}{n}$-oblivious transfer (OT) protocol, the sender $\mathcal{R}$ has private input $\mu = (\mu_1, \ldots, \mu_n) \subset M^n$ (and no private output) for some set M, while the chooser $\mathcal{I}$ has private input $\sigma \in [1,n]$ and private output μ_σ. The oblivious transfer (OT) protocol by Naor and Pinkas [NP01] guarantees information-theoretic privacy for $\mathcal{R}$, and computational privacy for $\mathcal{I}$. Intuitively, the in the Naor-Pinkas protocol, the sender oblivious-transfers one encryption key v_σ that is used to encrypt the actual database element μ_σ. The Naor and Pinkas [NP01] paper does not specify the encryption method, mentioning only that the encryption scheme must be semantically secure.

We propose to use Pedersen's commitment scheme instead of an encryption scheme. Let $K = (g; h)$ be the public key of the commitment scheme. The proposed variant of the Naor-Pinkas protocol works as follows:

1. $\mathcal{I}$ generates random $a, b \leftarrow \mathbb{Z}_q$ and sends $(A, B, C) \leftarrow (g^a, g^b, g^{ab-\sigma+1})$ to $\mathcal{R}$.
2. $\mathcal{R}$ performs the following, for $i \in [1,n]$: Generate random (r_i, s_i). Compute $w_i \leftarrow g^{r_i} A^{s_i}$, compute an encryption $y_i \leftarrow \mathsf{C}_K(\mu_i; v_i \mod q)$, where $v_i \leftarrow B^{r_i}(C \cdot g^{i-1})^{s_i}$. Send (w_i, y_i) to $\mathcal{I}$.
3. $\mathcal{I}$ computes $w_\sigma^b (= v_\sigma)$ and recovers $g^{\mu_\sigma} \leftarrow y_\sigma / h^{w_\sigma^b}$.

We denote this version of Naor-Pinkas protocol, where y_i is defined as $y_i = C_K(\mu_i, v_i)$, by $\binom{1}{n}$-$\mathsf{OT}_K(\mu;\sigma)$. As the end of this protocol, the verifier obtains commitments of all elements μ_i. Thus, the sender can argue in zero-knowledge for all $i \in [1,n]$ that the values μ_i satisfy some required conditions. We call such an OT protocol *verifiable*. (See [Lip03b] for a more precise definition.) $\mathcal{I}$ can "decrypt" y_σ with the "key" v_σ, given that the possible message space M is small enough for the exhaustive search on the set $\{g^x : x \in M\}$ to be practical. In the case of dichotomous RRT, $M = \{0,1\}$.

We define the sender privacy of an oblivious transfer protocol as follows. The chooser $\mathcal{I}^*$ chooses σ and two different vectors, $\mu[1] = (\mu[1]_1, \ldots, \mu[1]_n) \in M^n$ and $\mu[2] = (\mu[1]_1, \ldots, \mu[1]_n) \in M^n$, such that $\mu[1]_\sigma = \mu[2]_\sigma$. Denote an $\mathcal{I}^*$ that has

made such choices by $\mathcal{I}^*(\mu[1], \mu[2])$. He submits both tuples to the responder, who flips a fair coin $b \leftarrow_R [1, 2]$. After that, the chooser and the responder execute the protocol $\binom{1}{n}$-$\mathsf{OT}_K(\mu[b]; \sigma)$. After receiving $\mu[b]_\sigma$, $\mathcal{I}^*$ guesses the value of b. Let $\mathsf{Adv}_k^{\mathsf{lor}}(\mathcal{I}^*, \mathcal{R})$ be the probability that $\mathcal{I}^*$ guesses the correct b, where probability is taken over the internal coin tosses of $\mathcal{I}^*$ and $\mathcal{R}$. We say that the oblivious transfer protocol is *ε-sender-private*, if for any unbounded algorithm $\mathcal{I}^*$, $\mathsf{Adv}_k^{\mathsf{lor}}(\mathcal{I}^*, \mathcal{R}) \leq \varepsilon$.

Theorem 1. *Let $\binom{1}{n}$-$\mathsf{OT}_K(\cdot;\cdot)$ be the described oblivious transfer protocol. (a) If a malicious $\mathcal{R}^*$ can guess the value of σ with advantage ε, then he can solve the Decisional Diffie Hellman (DDH) problem with the same probability and in approximately the same time. (b) This protocol is $(m-d)(m-1)/q \leq m(m-1)/q$-sender-private, where $d := q \mod m$ and $m := |M|$.*

The security proof is omitted from this extended abstract due to the space constraints.

Zero-Knowledge Arguments. We will use zero-knowledge arguments (and not proofs) of knowledge in our protocol, since they are at the very least statistically hiding and computationally convincing. This property is important in a setting where a verifier must not be able to extract additional information even if he is given infinite time.

Our first protocol uses only two very standard statistical zero-knowledge arguments. The first one is an argument that a given value y_i (Pedersen-)commits to a Boolean value $\mu_i \in \{0, 1\}$. One can use standard disjunctive proofs for this. We denote the (possibly parallelized) argument that this holds for $i \in [1, n]$ by $\mathsf{AKEncBool}(y_1, \ldots, y_n)$. The second argument of knowledge, $\mathsf{AKLin}(y_1, \ldots, y_{n+1}; a, b)$, is an argument that the prover knows some set of values μ_i, for which y_i is a commitment of μ_i, and such that $\sum_{i \leq n} \mu_i + a\mu_{n+1} = b$. This argument of knowledge can be constructed from Pedersen's commitment scheme by computing $y \leftarrow \prod_{i \leq n} y_i \cdot y_{n+1}^a$ and then arguing that the result y is a commitment to b. Note that such an argument of knowledge is secure only when accompanied by zero-knowledge arguments of knowledge of the values μ_i; for this purpose, we employ $\mathsf{AKEncBool}(y_1, \ldots, y_{n+1})$ as described above.

5 Security Definitions

Next, we will give the definition of a weakly and strongly secure cryptographic RRT (CRRT). The security definitions will be in accordance with the ones in secure two-party computation. We will also explain why these requirements are relevant in the case of CRRT.

Assume we have a concrete variant of RRT, like RRT-W or RRT-IQ. Let Φ_p be the function that implements the desired functionality. For example, in the case of RRT-W, $\Phi_{p_{\mathsf{ct}}}(x)$ is a randomized function that with probability p_{ct} returns x, and with probability $1 - p_{\mathsf{ct}}$ returns $1 - x$. The ideal-world CRRT protocol, has three parties, the interviewer $\mathcal{I}$, the respondent $\mathcal{R}$, and the trusted third party $\mathcal{T}$. $\mathcal{R}$ has her type, $t_{\mathcal{R}}$ as her private input, while $\mathcal{I}$ has no private input. Then, $\mathcal{R}$ communicates $t_{\mathcal{R}}$ to $\mathcal{T}$, who selects the value $r_{\mathcal{R}} \leftarrow \Phi_{p_{\mathsf{ct}}}(t_{\mathcal{R}})$ and sends $r_{\mathcal{R}}$ to $\mathcal{I}$. After that, the private output of $\mathcal{I}$ will be $\Phi_{p_{\mathsf{ct}}}(t_{\mathcal{R}})$, while $\mathcal{R}$ will have no private output. It is required that at the end of the protocol, the participants will have no information about the private inputs and

outputs of their partners, except for what can be deduced from their own private inputs and outputs. In particular, $\mathcal{I}$ (resp. $\mathcal{R}$) has no information about the value of $t_{\mathcal{R}}$ (resp. $r_{\mathcal{R}}$), except what they can deduce from their private inputs and outputs.

In an ideal world, exactly the next three types of attacks are possible: a party can (a) refuse to participate in the protocol; (b) substitute his private input to the trusted third party with a different value; or (c) abort the protocol prematurely. In our case, the attack (c) is irrelevant, since $\mathcal{R}$ has no output. (Attack (c) models the case when the first party halts the protocol after receiving his private output but before the second party has enough information to compute her output.) Therefore, in an ideal-world RRT protocol, we cannot protect against a participant, who (a) refuses to participate in polling (*non-participation attack*) or (b) claims that her type is $1 - t_{\mathcal{R}}$, where $t_{\mathcal{R}}$ is her real type (*absolute denial attack*). No other attacks should be possible. Note that neither (a) nor (b) is traditionally considered an attack in the context of polling or voting. The argument here is game-theoretic, and the solutions must be proposed by mechanism design, instead of cryptography: namely, a non-manipulable mechanism (e.g., the algorithm with which the election winner is determined from all the collected votes) must be designed so that answering against one's true type (or non-participation) would not give more beneficial results to the respondent than the truthful answer.

On the other hand, as we stated, no other attacks should be allowed. This requirement is very strict, so we will explain why it is necessary in the RRT's context. Clearly, one must protect the privacy of $\mathcal{R}$, since this is the primarily goal of a RRT. It is also necessary to protect the privacy of $\mathcal{I}$, although the reason here is more subtle. Namely, if $\mathcal{R}$ obtains any additional information about $r_{\mathcal{R}}$ before the end of the protocol (for example, if she suspects that $r_{\mathcal{R}} \neq t_{\mathcal{R}}$), she might halt the protocol. Such a behavior by a malicious respondent might cause a bias in the poll, as already explained. (Halting the protocol while having no information on $r_{\mathcal{R}}$ is equivalent to the non-participation attack.) The third requirement on the protocol, of course, is that $\mathcal{I}$ either halts or receives $\Phi_{p_{\mathrm{ct}}}(x)$, where x is the input submitted by the $\mathcal{R}$.

In a real-world implementation, we want to replace $\mathcal{T}$ by a cryptographic protocol $\Pi = (\mathcal{R}, \mathcal{I})$ between $\mathcal{R}$ and $\mathcal{I}$. This protocol $(\mathcal{R}, \mathcal{I})$ is assumed to be "indistinguishable" from the ideal-world protocol, that is, with a high probability, it should be secure against all attacks that do not involve attacks (a) or (b). "Secure" means that the privacy of $\mathcal{R}$ (resp. $\mathcal{I}$) must be protected, if $\mathcal{R}$ (resp. $\mathcal{I}$) follows the protocol, and that $\mathcal{I}$ either halts, or receives the value $\Phi_{p_{\mathrm{ct}}}(x)$, where x was the submitted value of $\mathcal{R}$. The security of the respondent should be information-theoretical, while the security of interviewer can be computational. That is, a secure CRRT-W protocol must have the next three properties (here, k is the security parameter):

Privacy of Respondent: Let $\mathcal{I}^*$ be an algorithm. After the end of the protocol execution $(\mathcal{R}, \mathcal{I}^*)$, $\mathcal{I}^*$ will have no more information on $t_{\mathcal{R}}$ than it would have had after the execution of the ideal world protocol. That is, assuming that $\mathsf{view}_{\mathcal{I}^*}$ is his view of the protocol $(\mathcal{R}, \mathcal{I}^*)$, define $\mathsf{Adv}_k^{\mathsf{pri-r}}(\mathcal{R}, \mathcal{I}^*) := |\Pr[\mathcal{I}^*(\mathsf{view}_{\mathcal{I}^*}, r_{\mathcal{R}}) = t_{\mathcal{R}}] - \Pr[t_{\mathcal{R}} | r_{\mathcal{R}}]|$, where the probability is taken over the internal coin tosses of $\mathcal{I}^*$ and $\mathcal{R}$. We say that a CRRT protocol is *privacy-preserving for the respondent*, if $\mathsf{Adv}_k^{\mathsf{pri-r}}(\mathcal{R}, \mathcal{I}^*)$ is negligible (in k) for any unbounded adversary $\mathcal{I}^*$.

Privacy of Interviewer: Let $\mathcal{R}^*$ be an algorithm. Assume that $\mathcal{I}$ halts when $\mathcal{R}^*$ halts. After the end of the protocol execution $(\mathcal{R}^*, \mathcal{I})$, $\mathcal{R}^*$ will have no more information on $t_{\mathcal{R}}$ than it would have had after the execution of the ideal world protocol. That is, assuming that $\mathsf{view}_{\mathcal{R}^*}$ is her view of the protocol $(\mathcal{I}, \mathcal{R}^*)$, define $\mathsf{Adv}_k^{\mathsf{pri-i}}(\mathcal{R}^*, \mathcal{I}) := |\Pr[\mathcal{R}^*(\mathsf{view}_{\mathcal{R}^*}, t_{\mathcal{R}}) = r_{\mathcal{R}}] - \Pr[\mathcal{R}^*(t_{\mathcal{R}}) = r_{\mathcal{R}}]|$, where the probability is taken over the internal coin tosses of $\mathcal{R}^*$ and $\mathcal{I}$. We say that a CRRT protocol is *privacy-preserving for the interviewer*, if for any adversary $\mathcal{R}^*$, if $\mathsf{Adv}_k^{\mathsf{pri-i}}(\mathcal{R}^*, \mathcal{I}) \leq \varepsilon$ and $\mathcal{R}^*$ takes τ steps of computation then $\varepsilon\tau$ is negligible (in k).

Correctness: Let $\mathcal{R}^*(x)$ be an algorithm with private input x to the protocol $(\mathcal{R}^*, \mathcal{I})$. Assume that $\mathcal{I}$ halts when $\mathcal{R}^*$ halts. We require that at the end of the protocol execution $(\mathcal{R}^*, \mathcal{I})$, $\mathcal{I}$ will either halt, or otherwise receive $\Phi_{p_{\mathsf{ct}}}(x)$ with high probability. That is, assuming that $\mathsf{view}_{\mathcal{I}}$ is $\mathcal{I}$'s view of the protocol $(\mathcal{R}^*, \mathcal{I})$, define $\mathsf{Adv}_k^{\mathsf{crct}}(\mathcal{R}^*, \mathcal{I}) := 1 - \Pr[\mathcal{I}(\mathsf{view}_{\mathcal{I}}) = \Phi_{p_{\mathsf{ct}}}(x) | \mathcal{I} \text{ does not halt}]$, where the probability is taken over the internal coin tosses of $\mathcal{I}$ and $\mathcal{R}^*$. We say that a CRRT protocol is *correct*, if for any adversary $\mathcal{R}^*$, if $\mathsf{Adv}_k^{\mathsf{crct}}(\mathcal{R}^*, \mathcal{I}) = \varepsilon$ and $\mathcal{R}^*$ takes up to τ steps of computation then $\varepsilon\tau$ is negligible (in k).

We call a cryptographic RRT (CRRT) protocol *weakly secure* if it is privacy-preserving for the respondent and correct. We call CRRT protocol *(strongly) secure* if it is weakly secure and it is privacy-preserving for the interviewer. While a secure CRRT protocol is preferable in many situations, there are settings where a weakly secure CRRT protocol suffices, such as where halting can be easily detected and punished, or means for state recovery prevent modifications between a first and second attempt of executing the protocol.

6 Cryptographic RRT

We will propose three different CRRT-W protocols. In the first two protocols, the common parameters are $p_{\mathsf{ct}} = \ell/n > 1/2$ for $\ell, n \in \mathbb{Z}$; generators g and h whose mutual discrete logs are unknown (at least by $\mathcal{R}$); and $K = (g; h)$. $\mathcal{R}$ has private input $t = t_{\mathcal{R}}$, and $\mathcal{I}$'s private output is $r_{\mathcal{R}}$.

CRRT Protocol Based on Oblivious Transfer. Our first implementation of RRT-W is described in Protocol 1. The arguments of knowledge can be efficiently constructed, see Sect. 4. Here, we can use $\mathsf{AKLin}(y_1, \ldots, y_{n+1}; 2\ell - n; \ell)$ since $\sum_{i \leq n} \mu_i + (2\ell - n)\mu_{n+1} = \ell$ independently of the value of t. All the steps in this protocol must be authenticated.

If we take the number of bits that must be committed as the efficiency measure (communication complexity of the protocol), then our protocol has complexity $O(n)$. In the polling application, one can most probably assume that $n \leq 5$. The security proofs of this protocol follow directly from the properties of underlying primitives. As a direct corollary from Theorem 1, we get that Protocol 1 is privacy-preserving for respondent ($\mathsf{Adv}_k^{\mathsf{pri-r}}(\mathcal{R}, \mathcal{I}^*) \leq 2/q + O(1/q)$, where the constant comes in from the use of statistically-hiding zero-knowledge arguments). It is privacy preserving for interviewer, given the Decisional Diffie-Hellman (DDH) assumption. The correctness of

PRECOMPUTATION STEP:

1. $\mathcal{R}$ prepares n random bits $\mu_i \in \{0,1\}$ for $i \in [1,n]$, such that $\sum \mu_i = \ell$ if $t = 1$ and $\sum \mu_i = n - \ell$ if $t = 0$. Additionally, she sets $\mu_{n+1} \leftarrow 1 - t$.
2. $\mathcal{I}$ chooses an index $\sigma \in [1,n]$.

INTERACTIVE STEP:

1. $\mathcal{I}$ and $\mathcal{R}$ follow $\binom{1}{n}$-$\mathsf{OT}_K(g^{\mu_1}, \ldots, g^{\mu_n}; \sigma)$. $\mathcal{I}$ obtains g^{μ_σ}, and computes μ_σ from that.
2. $\mathcal{R}$ performs zero-knowledge arguments $\mathsf{AKEncBool}(y_1, \ldots, y_{n+1})$ and $\mathsf{AKLin}(y_1, \ldots, y_{n+1}; 2\ell - n; \ell)$ with $\mathcal{I}$ as the verifier.
3. $\mathcal{I}$ halts if the verification fails.

Protocol 1: A secure CRRT-W protocol based on oblivious transfer

this protocol follows from the properties of the zero-knowledge arguments used under the DDH assumption.

In a simplified weakly secure protocol based on the same idea, $\mathcal{R}$ commits to all μ_i by computing and publishing $y_i \leftarrow \mathsf{C}_K(\mu_i; \rho_i)$. Next, $\mathcal{R}$ argues that $\mathsf{AKEncBool}(y_1, \ldots, y_{n+1})$, and $\mathsf{AKLin}(y_1, \ldots, y_{n+1}; 2\ell - n; \ell)$. After that, $\mathcal{I}$ sends σ to $\mathcal{R}$, who then reveals μ_σ and ρ_σ. Upon obtaining these, $\mathcal{I}$ verifies the correctness of the previous corresponding commitment, outputting μ_σ.

CRRT from Coin-Flipping. Protocol 2 depicts a secure CRRT-W protocol with communication complexity $\Theta(d \log_2 n)$, where $d := \lceil 1/(1 - p_{\mathsf{ct}}) \rceil$, and $p_{\mathsf{ct}} = \ell/n$ as previously. While in the common RRT application one can usually assume that n is relatively small, this second protocol is useful in some specific game-theoretic applications where for the best outcome, p_{ct} must have a very specific value. The idea behind this protocol is that at least one of the integers $\mu + \nu + i\ell \mod n$ must be in interval $[0, \ell - 1]$, and at least one of them must be in interval $[\ell, n - 1]$. Hence, $\mathcal{I}$ gets necessary proofs for both the 0 and the 1 answer, which is sufficient for his goal. For his choice to be accepted, he must accompany the corresponding r with $\mathcal{R}$-s signature on his commitment on σ.

PRECOMPUTATION STEP:

1. $\mathcal{R}$ chooses a random $\mu \leftarrow_R [0, n-1]$.
2. $\mathcal{I}$ chooses random $\nu \leftarrow_R [0, n-1]$ and $\sigma \leftarrow_R [0, d-1]$.

INTERACTIVE STEP:

1. $\mathcal{R}$ commits to t and μ, and sends the commitments to $\mathcal{I}$.
2. $\mathcal{I}$ commits to σ, by setting $y \leftarrow \mathsf{C}_K(\sigma; \rho)$ for some random ρ. He sends ν and y to $\mathcal{R}$, together with a zero-knowledge argument that y is a commitment of some $i \in [0, d-1]$.
3. $\mathcal{R}$ verifies the argument. She computes values μ'_i, for $i \in [0, d-1]$, such that $\mu'_i = t \iff (\mu + \nu + i\ell \mod n) < \ell$. She signs y, and sends her signature together with $\{\mu'_i\}$ and the next zero-knowledge argument for every $i \in [0, d-1]$: $[\mu'_i = t \iff (\mu + \nu + i\ell \mod n) < \ell]$.
4. After that, $\mathcal{I}$ sets $r_{\mathcal{R}} \leftarrow \mu'_\sigma$. He will accompany this with $\mathcal{R}$-s signature on the commitment, so that both $\mathcal{R}$ and third parties can verify it.

Protocol 2: A secure CRRT-W protocol based on coin-flipping

PRECOMPUTATION STEP:

1. $\mathcal{I}$ chooses random $u_0 \leftarrow_R [0,1]$, $u_1 \leftarrow_R [0,1]$. He generates quantum states $|\psi_0\rangle = \sqrt{p_{\mathsf{ct}}}|u_0\rangle + \sqrt{1-p_{\mathsf{ct}}}|1-u_0\rangle$, $|\psi_1\rangle = \sqrt{p_{\mathsf{ct}}}|u_1\rangle + \sqrt{1-p_{\mathsf{ct}}}|1-u_1\rangle$.
2. $\mathcal{R}$ chooses a random $i \leftarrow_R [0,1]$.

INTERACTIVE STEP:

1. $\mathcal{I}$ sends $|\psi_0\rangle$ and $|\psi_1\rangle$ to $\mathcal{R}$.
2. $\mathcal{R}$ sends i to $\mathcal{I}$.
3. $\mathcal{I}$ sends u_i to $\mathcal{R}$.
4. $\mathcal{R}$ measures the state $|\psi_i\rangle$ in the basis $|\psi_{u_i}\rangle = \sqrt{p_{\mathsf{ct}}}|u_i\rangle + \sqrt{1-p_{\mathsf{ct}}}|1-u_i\rangle$, $|\psi_{u_i}^{\perp}\rangle = \sqrt{1-p_{\mathsf{ct}}}|u_i\rangle - \sqrt{p_{\mathsf{ct}}}|1-u_i\rangle$ and halts if the result is not $|\psi_{u_i}\rangle$.
5. If the verification is passed, $\mathcal{R}$ performs the transformation $|0\rangle \rightarrow |t\rangle$, $|1\rangle \rightarrow |1-t\rangle$ on the state $|\psi_{1-i}\rangle$ and sends it back to $\mathcal{I}$.
6. $\mathcal{I}$ measures the state in the basis $|0\rangle$, $|1\rangle$, gets outcome s. $\mathcal{I}$ outputs $r \leftarrow u_i \oplus s$.

Protocol 3: A quantum CRRT-W protocol.

A weakly secure version of this protocol is especially efficient. There, one should set $d \leftarrow 1$, and omit the steps in Protocol 2 that depend on σ being greater than 1. (E.g., there is no need to commit to σ anymore.) Thus, such a protocol would have communication complexity $\Theta(\log_2 n)$. Now, $p_{\mathsf{ct}} > 1/2$ (otherwise one could just do a bit-flip on the answers), and hence $d > 2$. On the other hand, the privacy of respondents is in danger if say $p_{\mathsf{ct}} \geq 3/4$. Thus, we may assume that $d \in [3,4]$. Therefore, Protocol 2 will be more communication-efficient than Protocol 1 as soon as $n/\log_2 n > 4 \geq d$, or $n \geq 16$. The weakly secure version will be *always* more communication-efficient.

This protocol is especially efficient if the used commitment scheme is an integer commitment scheme. In this case, to argue that $(\mu + \nu + i\ell \mod n) < \ell$ one only must do the next two simple steps: first, argue that $\mu + \nu + i\ell = z + en$ for some z, e, and then, argue that $z \in [0, \ell - 1]$. This can be done efficiently by using the range proofs from [Lip03a]. One can also use Pedersen's scheme, but this would result in more complicated arguments.

Quantum-Cryptographic RRT. The next *quantum CRRT protocol* (see Protocol 3) works also for irrational p_{ct}, and provides a relaxed form of information-theoretic security to *both* parties. While not secure by our previous definitions, it provides meaningfully low bounds on the probabilities of success for a cheater. Namely, (a) if dishonest, $\mathcal{R}$ cannot make his vote count as more than $\sqrt{2}$ votes: if $p_{\mathsf{ct}} = \frac{1}{2} + \varepsilon$, then $p_{\mathsf{adv}} \leq \frac{1}{2} + \sqrt{2}\varepsilon$ (The full version of this paper has a slightly better bound with a more complicated expression for p_{adv}). (b) if dishonest strategy allows $\mathcal{I}$ to learn t with probability $p_{\mathsf{ct}} + \varepsilon$, it also leads to $\mathcal{I}$ being caught cheating with probability at least $\frac{2p_{\mathsf{ct}}-1}{2}\varepsilon$. This form of security (information-theoretic security with relaxed definitions) is common for quantum protocols for tasks like bit commitment or coin flipping. The security guarantees of our quantum protocol compare quite well to ones achieved for those tasks. A desirable property of this quantum protocol is that it can be implemented by using contemporary technology, since it only involves transmitting and measuring single qubits, and no maintaining of coherent multi-qubit states.

To show the main ideas behind quantum protocol, we now show how to analyze a simplified version of protocol 3. The security proof for the full protocol is quite complicated and will be given in the full version of this paper.

The simplified version of Protocol 3 is: (1) $\mathcal{I}$ chooses a random $u \leftarrow_R [0,1]$, prepares a quantum bit in the state $|\psi_u\rangle = \sqrt{p_{\mathsf{ct}}}|u\rangle + \sqrt{1-p_{\mathsf{ct}}}|1-u\rangle$ and sends it to $\mathcal{R}$; (2) $\mathcal{R}$ performs a bit flip if her type $t = 1$, and sends the quantum bit back to $\mathcal{I}$; (3) $\mathcal{I}$ measures the state in the computational basis $|0\rangle$, $|1\rangle$, gets answer s. The answer is $r = u \oplus s$. If both parties are honest, the state returned by respondent is unchanged: $\sqrt{p_{\mathsf{ct}}}|u\rangle + \sqrt{1-p_{\mathsf{ct}}}|1-u\rangle$ if $t = 0$ and $\sqrt{p_{\mathsf{ct}}}|1-u\rangle + \sqrt{1-p_{\mathsf{ct}}}|u\rangle$ if $t = 1$. Measuring this state gives the correct answer with probability $1 - p_{\mathsf{ct}}$. Next, we show that respondent is unable to misuse this protocol.

Theorem 2. *For any respondent's strategy $\mathcal{R}^*$, the probability of honest interviewer $\mathcal{I}$ getting $r = 1$ is between $1-p_{\mathsf{ct}}$ and p_{ct}. Therefore, the previous protocol is both correct and privacy-preserving for the interviewer.*

Proof. We show that the probability of $r = 1$ is at most p_{ct}. The other direction is similar. We first modify the (simplified) protocol by making $\mathcal{R}^*$ to measure the state and send the measured result to $\mathcal{I}$, this does not change the result of the honest protocol since the measurement remains the same. Also, any cheating strategy for $\mathcal{R}^*$ in the original protocol can be used in the new protocol as well. So, it is sufficient to bound the probability of $r = 1$ in the new protocol. The answer is $r = 1$ if $\mathcal{I}$ sent $|\psi_i\rangle$ and $\mathcal{R}^*$ sends back j, with $i = j$. By a well-known fact, the maximum success probability with what one can distinguish two qubits is $1/2 + \sin\beta/2$, where β is the angle between two qubits. The rest is a calculation: to determine the angle β between $|\psi_0\rangle$ and $|\psi_1\rangle$, it suffices to determine the inner product which is $\sin\beta = 2\sqrt{p_{\mathsf{ct}}(1-p_{\mathsf{ct}})}$. Therefore, $\cos\beta = \sqrt{1-\sin^2\beta} = 2p_{\mathsf{ct}} - 1$ and $\frac{1}{2} + \frac{\cos\beta}{2} = p_{\mathsf{ct}}$. □

On the other hand, when using this simplified version, a dishonest interviewer $\mathcal{I}^*$ can always learn t with probability 1. Namely, it suffices to send the state $|0\rangle$. If $t = 0$, $\mathcal{R}$ sends $|0\rangle$ back unchanged. If $t = 1$, $\mathcal{R}$ applies a bit flip. The state becomes $|1\rangle$. $\mathcal{I}$ can then distinguish $|0\rangle$ from $|1\rangle$ with certainty by a measurement in the computational basis.

Note that this is similar to a classical "protocol", where $\mathcal{I}$ first generates a random u and sends a bit i that is equal to u with probability p_{ct} and $1-u$ with probability $1-p_{\mathsf{ct}}$. $\mathcal{R}$ then flips the bit if $t = 1$ and sends it back unchanged if $t = 0$. The interviewer XORs it with u, getting t with probability p_{ct} and $1 - t$ with probability $1 - p_{\mathsf{ct}}$. In this "protocol", $\mathcal{R}$ can never cheat. However, $\mathcal{I}^*$ can learn t with probability 1 by just remembering i and XORing the answer with i instead of u. In the classical world, this flaw is fatal because $\mathcal{I}$ cannot prove that he has generated i from the correct probability distribution and has not kept a copy of i for himself. In the quantum case, $\mathcal{I}$ can prove to $\mathcal{R}$ that he has correctly prepared the quantum state. Then, we get Protocol 3 with $\mathcal{I}$ sending two states $|\psi_{u_0}\rangle$ and $|\psi_{u_1}\rangle$, one of which is verified and the other is used for transmitting t. A detailed analysis of this protocol is omitted from this extended abstract.

7 Protocols for Other RRTs and Extensions

Protocol for Cryptographic RRT-IQ. Recall that in one version of RRT-IQ, the respondent would reply with his true opinion $t_{\mathcal{R}}$ with a rational probability $p_{\mathsf{ct}} = \ell/n$, while he would otherwise flip a coin and answer whether it came up tails. Like for CRRT-W, it is important to guarantee the use of correct distributions. Protocol 1 can be easily changed to work for this version of RRT-IQ. Instead of n random bits, $\mathcal{R}$ prepares $2n$ random bits μ_i, so that either $\sum_{i=1}^{n} \mu_i \in \{\ell, n-\ell\}$ and $\sum_{i=n+1}^{2n} \mu_i = n/2$ or $\sum_{i=n+1}^{2n} \mu_i \in \{\ell, n-\ell\}$ and $\sum_{i=1}^{n} \mu_i = n/2$. She then uses standard techniques to prove that the bits were prepared correctly, after which $\mathcal{I}$ chooses one of the $2n$ bits by using the verifiable oblivious transfer protocol. (Here, of course, n must be even.)

Protocol for Cryptographic PRRT-BD. The next protocol is a modification of Protocol 1 as well. Let p_i be such that $p_{\mathsf{ct}} + \sum_{i \in [1,m]} p_i = 1$, and assume that every respondent has a type $t_{\mathcal{R}} \in [1, m]$. Assume $p_{\mathsf{ct}} = \ell/n$, $p_i = \ell_i/n$ and that $p_i = 0$ if $i \notin [1, m]$. Assume $D \geq \max(\ell, \ell_1, \ldots, \ell_m) + 1$. The respondent prepares n numbers D^{μ_i}, such that $\sharp\{i : \mu_i = t_{\mathcal{R}}\} = \ell_{t_{\mathcal{R}}} + \ell$, and $\sharp\{i : \mu_i = j\} = \ell_j$, if $j \neq t_{\mathcal{R}}$. Then the interviewer and respondent will execute a variant of OT with choice σ, during which the interviewer only gets to know the value μ_σ. Then the respondent argues that the sum of all commitments is a commitment to the value $\sum \ell_i D^{\mu_i} + \ell D^j$, for some $j \in [1, m]$, by using range-proofs in exponents [LAN02]. (A more efficient proof methodology is available when D is a prime [LAN02], given that one uses an integer commitment scheme.) Additionally, she argues that every single commitment corresponds to a value D^i for $i \in [1, m]$, also using range-proofs of exponents [LAN02]. After the OT step, the interviewer gets g^{μ_σ}, and recovers μ_σ from it efficiently. (Note that $m \leq 10$ is typical in the context of polling.)

Extensions to Hierarchies of Interviewers. One can consider a hierarchy of interviewers, reporting to some central authority. If there is a trust relationship between these two types of parties, no changes to our protocol would be required. However, if the central authority would like to be able to avoid having to trust interviewers, the following modifications could be performed. First, each respondent would have to authenticate the transcript he generates, whether with a standard signature scheme, a group signature scheme, etc. Second, and in order to prevent collusions between interviewers and respondents, the interviewers must not be allowed to know the choice σ made in a particular interview. Thus, the triple (A, B, C) normally generated by the interviewer during the Naor-Pinkas OT protocol would instead have to be generated by the central authority, and kept secret by the same. More efficient versions of *proxy* OT satisfying our other requirements are beneficial for this application.

Full version. Due to the space constraints, we had to omit the security proof of the new verifiable oblivious transfer protocol and a detailed analysis of the quantum RRT protocol. The full version of this paper is available from the IACR eprint archive.

Acknowledgments. The third author was supported by the Finnish Defense Forces Research Institute of Technology and by the Finnish Academy of Sciences. We would like to thank Jouni K. Seppänen and Benny Pinkas for useful comments.

References

[BR99] Mihir Bellare and Ronald Rivest. Translucent Cryptography — An Alternative to Key Escrow, and Its Implementation via Fractional Oblivious Transfer. *Journal of Cryptology*, 12(2):117–139, 1999.

[CK88] Claude Crépeau and Joe Kilian. Achieving Oblivious Transfer Using Weakened Security Assumptions (Extended Abstract). In *29th Annual Symposium on Foundations of Computer Science*, pages 42–52, White Plains, New York, USA, October24–26 1988. IEEE Computer Society Press.

[CM88] Arijit Chaudhuri and Rahul Mukerjee. *Randomized Response: Theory and Techniques*, volume 95 of *Statistics: Textbooks and Monographs*. Marcel Dekker, Inc., 1988. ISBN: 0824777859.

[Cré97] Claude Crépeau. Efficient Cryptographic Protocols Based on Noisy Channels. In Walter Fumy, editor, *Advances in Cryptology — EUROCRYPT '97*, volume 1233 of *Lecture Notes in Computer Science*, pages 306–317, Konstanz, Germany, 11–15 May 1997. Springer-Verlag.

[GASH69] Bernard G. Greenberg, Abdel-Latif A. Abul-Ela, Walt R. Simmons, and Daniel G. Horvitz. The Unrelated Question Randomized Response Model: Theoretical Framework. *Journal of the American Statistical Association*, 64(326):520–539, June 1969.

[KANG99] Hiroaki Kikuchi, Jin Akiyama, Gisaku Nakamura, and Howard Gobioff. Stochastic Voting Protocol To Protect Voters Privacy. In *1999 IEEE Workshop on Internet Applications*, pages 103–111, July 26–27 1999.

[Lai03] Chi Sung Laih, editor. *Advances on Cryptology — ASIACRYPT 2003*, volume 2894 of *Lecture Notes in Computer Science*, Taipei, Taiwan, November 30–December 4 2003. Springer-Verlag.

[LAN02] Helger Lipmaa, N. Asokan, and Valtteri Niemi. Secure Vickrey Auctions without Threshold Trust. In Matt Blaze, editor, *Financial Cryptography — Sixth International Conference*, volume 2357 of *Lecture Notes in Computer Science*, pages 87–101, Southhampton Beach, Bermuda, March 11–14 2002. Springer-Verlag.

[Lip03a] Helger Lipmaa. On Diophantine Complexity and Statistical Zero-Knowledge Arguments. In Laih [Lai03], pages 398–415.

[Lip03b] Helger Lipmaa. Verifiable Homomorphic Oblivious Transfer and Private Equality Test. In Laih [Lai03], pages 416–433.

[NP01] Moni Naor and Benny Pinkas. Efficient Oblivious Transfer Protocols. In *Proceedings of the Twelfth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 448–457, Washington, DC, USA, January 7–9 2001. ACM Press.

[Ped91] Torben P. Pedersen. Non-Interactive and Information-Theoretic Secure Verifiable Secret Sharing. In J. Feigenbaum, editor, *Advances in Cryptology—CRYPTO '91*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140, Santa Barbara, California, USA, August 11–15 1991. Springer-Verlag, 1992.

[War65] Stanley L. Warner. Randomized Response: A Survey Technique for Eliminating Evasive Answer Bias. *Journal of the American Statistical Association*, 60(309):63–69, March 1965.