Noncryptographic Selection Protocols
(Extended Abstract)

Uriel Feige
Weizmann Institute, Rehovot, Israel
feige@wisdom.weizmann.ac.il

October 24, 2000

Abstract

Selection tasks generalize some well studied problems, such as collective coin flipping and leader election. We present new selection protocols in the full information model, and new negative results. In particular, when there are $(1+\delta)n/2$ good players, we show a protocol that chooses a good leader with probability $\Omega(\delta^{1.65})$, and show that every leader election protocol has success probability $O(\delta^{1-\epsilon})$, for every $\epsilon > 0$. Previously known protocols for this problem have success probability that is exponentially small in $1/\delta$, and no nontrivial upper bounds on the success probability were known.

1 Introduction

In the full information model there are n players and all communication is through broadcast operations. It is assumed that broadcasting is a reliable atomic act: all other players receive the message uncorrupted together with the identity of the player who broadcasts the message. The players need to jointly perform some task (such as elect a leader). Some of the players may be faulty. The good players do not know which of the other players is faulty. The faulty players may use unlimited computational powers and coordinate their actions so as to make the outcome of the task favorable to them (e.g., to have a faulty player elected as the leader). This is modeled by a computationally unlimited adversary who picks the set of faulty players before the execution of the protocol begins, and thereafter coordinates their actions. It is assumed that players can individually generate random bits (each player has his own private coin), and that values of future coin tosses of nonfaulty players are unpredictable, unbiased and independent of any other event. (The coins of faulty players cannot be trusted.) The full information model does not support standard cryptographic protocols (which require some assumption such as private communication channels, or computational intractability).

The problem of designing a collective coin flipping protocol in this model was suggested by Ben-Or and Linial [4]. The task of the players is to come up with a common random bit. A protocol for this task is *resilient* if the bit is somewhat random – regardless of the behavior of the bad players, there is some probability bounded away from 0 with which the bit receives each of its two possible values. Naturally, the good players would need to follow a randomized algorithm. A case that received special attention is one round protocols in which the value of the global coin is obtained by each player supplying one bit of input to some n-argument predetermined Boolean function. Good players supply random bits, whereas faulty players supply arbitrary bits that may depend on the bits of the good players. Ajtai and Linial [2] designed a function that is resilient whenever the number of bad players is smaller than $n/\log^2 n$, and Kahn, Kalai and Linial [9] showed that no function can be resilient against more than $n/\log n$ players.

Leader election is the task of selecting one player out of the n. A protocol for leader election is resilient if the probability of choosing a good leader is bounded away from 0. Leader election protocols can be used as protocols for collective coin flipping, by having the leader toss the coin.

Saks [14] showed that the "Baton Passing" game, in which each player receiving the baton passes it to a random player who did not yet have it, and the last player left with the baton is the leader, is resilient when the number of bad players is below $n/\log n$. (The bad players of course try to always pass the baton to a good player.) The Baton Passing game takes $n-1$ rounds. Saks also observed that no leader election protocol can be resilient if half the players are bad (a proof appears in [6]).

The Baton Passing game was modified by Alon and Naor [3] to a game in which in each round a player can pass the Baton to one of two players, where these two players are determined by the complete history of the protocol. They showed that the modified game is resilient against $(1-\epsilon)n/3$ bad players. Boppana and Narayanan [6] improved the analysis of such games and shown them to be resilient against $(1-\epsilon)n/2$ bad players, for every $\epsilon > 0$. It is not known whether the modified baton passing game can be implemented when the computation time of players is polynomial in n, due to the complexity of figuring out which are the two players to which a player is allowed to pass the Baton.

The question of reducing the number of rounds in leader election protocols (and hence also coin flipping protocols) was studied in [7, 11, 16] and by Russell and Zuckerman [13] who designed a polynomial time computable leader election protocol resilient against $(1-\epsilon)n/2$ bad players that takes only $\log^* n + O(1)$ rounds. Russell, Saks and Zuckerman [12] show that $\Omega(\log^* n)$ are necessary for resilient leader election if in every round the good players each sends one unbiased random bit. (Neither the results of [12] nor those of [9] apply when good players are allowed to send biased bits.)

1.1 Marginal majorities

In the current paper we study the case in which there is only a slight majority of good players. The number of good players is represented as $(1 + \delta)n/2$, for some small $\delta > 0$. Let $p(n, \delta)$ denote the probability of choosing a good leader in this case, under an optimal leader election protocol. Then previous work ([6], following [3]) established:

Theorem 1 *For every $\delta > 0$ there exists $\epsilon > 0$ such that for every n, $p(n, \delta) > \epsilon$.*

Theorem 1 establishes that $p(n, \delta)$ can be bounded from below as a function of δ, independently of n. Hence we shall sometimes omit n from the notation and use $p(\delta)$. We are interested in obtaining the best possible bounds on $p(\delta)$. The proof in [6] gives a lower bound that is a little worse than exponentially small, namely $p(\delta) \geq c^{-(\log 1/\delta)^2/\delta}$ [5], for some $c > 1$. An alternative proof of Theorem 1 is implicit in [13], who show a protocol that has constant probability of reducing the number of players to some constant n_0 that depends on δ, while maintaining a majority of good players. Thereafter, it is suggested to use the protocol of [6] on the remaining n_0 players, but of course at this stage, any other protocol with positive success probability would insure that $p(n, \delta)$ will be lower bounded in terms of δ alone. The question of whether there is a more favorable protocol to use at this point was not addressed in [13].

In terms of upper bounds, trivially $p(n, \delta) \leq (1 + \delta)/2$, as the faulty players can just play honestly. The author is not aware of any other published upper bound on $p(\delta)$.

For collective coin flipping and δ as above, let $r(\delta)$ be the minimum of the two probabilities that the coin comes up 1 or 0, under the worst case strategy for the adversary in the best coin flipping protocol. Letting the leader flip the coin we obtain $r(\delta) \geq p(\delta)/2$. An upper bound on $r(\delta)$ follows from the following theorem of [4, 8].

Theorem 2 *Let p_v be the probability that the outcome of an n player full information protocol is v if all players play at random. Then for every $1 \leq t \leq n$, there is a set of t "influential" players who have a strategy under which with probability at least $(p_v)^{1-t/n}$ the outcome is v.*

Theorem 2 gives a nontrivial upper bound for collective coin flipping. When all players play randomly, then w.l.o.g. we can assume that p_0, the probability that the coin comes up 0, is at least $1/2$. Hence there is a set of $(1 - \delta)n/2$ influential players that can force the coin to 0 with probability at least $(1/2)^{1-(1-\delta)/2}$. Hence $r(\delta) \leq 1 - (1/2)^{(1+\delta)/2}$, which is less than 0.293 when δ is small enough. Note that this upper bound remains bounded away from 0 when δ tends to 0.

Theorem 2 is not directly applicable to leader election protocols, because there the choices of which are the bad players and what is a bad outcome are correlated. For example, there is a protocol that elects a good leader with probability $(n - 1)/n$ when at most one player is bad. See Section C.1 in the appendix.

1.2 Our results

We present a simple protocol, based on the *lightest bin* principle, that can be used to reduce the number of players while maintaining (in a probabilistic sense) the fraction of good players.

This protocol and simple variations on it have several immediate consequences:

- It gives a simplified proof of Theorem 1, that $p(n, \delta)$ can be lower bounded by a function of δ, independent of n.
- It gives a simplified proof of the result of [13] that for every fixed δ, leader election can be performed in $\log^* n + O(1)$ rounds.
- It gives a quasipolynomial lower bound on $p(\delta)$, namely $p(\delta) \geq \delta^{O(\log 1/\delta)}$. The protocol achieving this depends only on n but not on δ, so this bound can be achieved for all $\delta \geq 1/n$ simultaneously.
- It shows that leader election and collective coin flipping can be performed essentially with the same success probability. Namely, $p(\delta) = \Theta(r(\delta))$.

The main technical contribution of the paper is the proof of the following theorem:

Theorem 3 *There are universal constants $c_1 \geq c_2 > 0$ such that:*

1. *For every δ, there is a leader election protocol (depending on δ and n) achieving $p(\delta) \geq \Omega(\delta^{c_1})$.*
2. *For every leader election protocol, $p(\delta) \leq O(\delta^{c_2})$*

The lower bound is proved by establishing a connection between leader election protocols and monotone circuits for majority. $p(\delta)$ can be lower bounded as a function of the depth of such circuits. The upper bound is proved by representing leader election protocols (or rather, collective coin flipping protocols) as Markov chains, and using martingale tail inequalities. The values that our proofs give are $c_1 < 1.65$ and $c_2 > 1 - \epsilon$ for every $\epsilon > 0$.

1.3 Selection versus sampling

Collective coin flipping can be viewed as a multiplayer random sampling problem, in which n players want to sample at random from $\{0, 1\}$. Leader election can also be viewed as a sampling problem, in which the n players wish to sample at random from $\{1 \ldots, n\}$, making the sampled player the leader.

In establishing our results, we found it useful to study leader election problems in a somewhat different framework, that we call selection protocols. In these protocols, a set of players need to jointly decide on an action (such as selecting a leader, or flipping a coin). Some of the players may be faulty, and some of the actions may be bad. The goal is to design protocols that allow good players to force the choice of a good action, provided there are sufficiently many good players and good actions. We call such protocols robust. We found this view as useful both for designing sampling protocols (by letting the good players just play at random in a robust protocol) and in proving negative results (similar to the impossibility of leader election with a majority of faulty players). The more general framework is presented in Section 2.

We note that Goldreich, Goldwasser and Linial [8] presented a different generalized framework for fault tolerant computation in the full information model. In their framework, each player has an input (unknown to other players), and all players jointly compute the value of a known n-argument function f on these inputs. The faulty players may influence the outcome of the computation towards a value v that they favor, by modifying their inputs, based on what they learn on the inputs of the good players. Goldreich et.al. study the influence the bad players have when the inputs to the good players are chosen at random. They give protocols that in some settings limit the influence of the bad players to be within a constant multiplicative factor of the bounds given in Theorem 2.

1.4 Roadmap

In Section 2 we describe our general framework for selection protocols. Other sections can be understood without reading this section. Section 3 presents a simple protocol, perhaps even practical, for tasks such as leader election. The success probability of this protocol is improved in Section 4, by adding a second phase to the protocol. Our upper bound on the success probability of collective coin flipping protocols is presented in Section 5. The appendix contains the technical part of the proof of the upper bound (Section A), upper bounds for specialized problems (Section B), and toy examples of protocols (Section C).

2 A general framework for selection protocols

There are n players, some of which are faulty. There are m possible candidates, some of which are bad. It is not known a-priori which are the faulty players and which are the bad candidates. We are interested in protocols in which the players collectively select a candidate. Our protocols are robust in the following sense. If at least a *coalition* of players is nonfaulty and a *quorum* of candidates is good (we shall define coalitions and quorums shortly), then the nonfaulty players can force the outcome of the protocol to be a good candidate. We now describe our setting formally.

There is a set $\{P_1, \ldots, P_n\}$ of n players and a collection $\mathcal{S} = \{S_1, S_2, \ldots\}$ of *coalitions*, where each coalition is a subset of the players. There is a set $\{a_1, \ldots a_m\}$ of m possible candidates, and a collection $\mathcal{Q} = \{Q_1, Q_2, \ldots\}$ of *quorums*, where each quorum is a subset of candidates. Players communicate by broadcasting messages from a fixed alphabet. A *protocol* specifies the order in which players speak, and the candidate selected when the protocol ends (depending on the actual messages broadcast by the players). A *strategy* for a player tells the player which character to broadcast as a function of the complete history of broadcasts seen so far. A *randomized strategy* is a probability distribution over strategies.

Definition 1 *A protocol is* robust *with respect to $\mathcal{Q}$ and $\mathcal{S}$ if for every coalition $S \in \mathcal{S}$ and for every quorum $Q \in \mathcal{Q}$, the players in S have a strategy such that regardless of the strategy of the other players, the candidate a eventually selected satisfies $a \in Q$.*

A pair $(\mathcal{S}, \mathcal{Q})$ specifies a selection problem. We are interested in characterizing the selection problems for which a robust protocol exists.

Definition 2 *Let j be a positive integer. Collection $\mathcal{S}$ is j-*intersecting *if there are $j-1$ mutually disjoint coalitions, but no j coalitions are mutually disjoint.*

As an example of a j-intersecting collection, let $\mathcal{S}$ contain all subsets of cardinality greater than n/j.

Definition 3 *Let i be a positive integer. Collection $\mathcal{Q}$ is an i-*quorum system *if every i quorums have a common intersection, and there are $i+1$ quorums that do not have a common intersection.*

As an example of an i-quorum system, let $\mathcal{Q}$ contain all subsets of cardinality greater than $m(1 - 1/i)$.

Theorem 4 *Let $\mathcal{S}$ be j-intersecting and $\mathcal{Q}$ be an i-quorum. Then there is a robust protocol with respect to $\mathcal{Q}$ and $\mathcal{S}$ if and only if $i \geq j - 1$.*

Proof: Assume $i \geq j - 1$, and consider the following protocol. Each player announces an arbitrary quorum $Q \in \mathcal{Q}$. Call a quorum popular if all players from some coalition announced this quorum. The candidate a chosen by the protocol is the first candidate that belongs to every popular quorum. There must be such a candidate a, because there can be at most $j-1$ different popular quorums, and these $j-1$ quorums have a common intersection. The protocol is robust because all players in a coalition S can announce the same quorum Q. This completes the *if* direction.

Assume $i < j-1$. There are $j-1$ coalitions in $\mathcal{S}$ that are mutually disjoint. W.l.o.g., let them be $S_1, \ldots, S_{j-1}$. There are $i+1$ quorums in $\mathcal{Q}$ that do not have a common intersection.

W.l.o.g., let them be $Q_1, \ldots, Q_{i+1}$. Then it cannot be that for every $k \leq i+1 \leq j$ there is strategy for the players in S_k to force the candidate to be chosen from Q_k, because there is no candidate in the intersection of the Q_ks. This completes the *only if* direction. □

Theorem 4 gives a complete characterization of selection problems that have robust protocols in the full information model. Robust protocols are characterized by an existential statement: the coalition of good players has a strategy that forces the candidate to be selected from the good quorum. However, we will be interested in cases when the good players do not know which is the good coalition, and which is the good quorum. Hence the good players might not follow the favorable strategy. We would like a randomized strategy for the good players that maximizes the probability of choosing a good candidate, regardless of which is the good coalition, which is the good quorum, and the strategy employed by the bad players. Here we may assume that prior to the beginning of the execution of the protocol, an adversary makes some of the players faulty and some of the candidates bad, but leaves at least one nonfaulty coalition and at least one good quorum. During the execution of the protocol, the adversary has full control of the messages sent by the faulty players.

Corollary 5 *Let $\mathcal{S}$ be j-intersecting and $\mathcal{Q}$ be an i-quorum, with $i \geq j-1$. Let $|\mathcal{Q}|$ denote the number of quorums, and let S be a minimal coalition of maximum cardinality (minimal in the sense that it does not properly contain another coalition). Then there protocols and randomized strategies that have success probability at least $|\mathcal{Q}|^{-|S|}$, regardless of the strategy of the adversary. Furthermore, if $j = 2$, then there are protocols and randomized strategies with success probability at least $m^{-|S|}$.*

Proof: For the general case, use the robust protocol of Theorem 4, and the randomized strategy in which each player chooses a random quorum. For the case $j = 2$, let each player choose a random candidate, and select a candidate chosen by a coalition (if there is none, select an arbitrary candidate). □

3 The *lightest bin* protocol

In this section we present a simple protocol for the committee election problem. In this problem, there are n players, at least k of which are good, and they want to elect a good committee of c players, where a committee is good if it contains at least one good player. This problem is solvable if and only if $k > n/(c+1)$ (Theorem 4). The case $c = 1$ corresponds to leader election, but we shall also find the case $c = 2$ very useful. The case $c \simeq \log n$ was suggested by Moni Naor (private communication) as having potential cryptographic applications.

We use the following notation and conventions:

S – Set of all players.

X – Temporary set created during the protocol.
L – Final set, outcome of the protocol.
n – Number of players.
k – Number of good players.
c – Size of the committee to be chosen.
δ – The *advantage*, $\delta = \frac{k(c+1)}{n} - 1$.
$p(n, k, c)$ – Probability of *success* – that of choosing a good committee.

Our protocols return a set $L \subset S$. If $|L| < c$, we add players to L arbitrarily.

As noted earlier, we need $\delta > 0$ for the committee election problem to be solvable, which in fact implies $\delta \geq 1/n$. We shall use the notation $p(\delta, c)$ if all we assume on n and k is that $k(c+1)/n \geq 1 + \delta$, and omit c from the notation when $c = 1$.

Our protocols can be described as games of throwing balls into bins. The game proceeds in rounds. There are several bins, and each player gets to throw his ball into a random bin. The bin that then contains the smallest number of balls is called the *lightest bin*. The players who have their balls in the lightest bin continue to the next round, and all other players are discarded. The balls are returned to the players, and the protocol is repeated recursively. When the lightest bin contains not more than c players, these players become the elected committee.

Of course, the bad players need not throw their balls into random bins. Rather, they wait to first see where the good player's balls land, and then try to place as many of their own balls as possible in the lightest bin. However, any bin that contains many bad balls will not be light, and will not continue to the next round. Hence even though the number of players is reduced in each round, the proportion of good players remains favorable.

The simplest version of our protocol works for leader election when k, the number of good players, is an exact power of 2, and $n = 2k - 1$.

Lightest Bin Protocol (simplified version):

1. $X \leftarrow S$.
2. Repeat while $|X| > 1$:
   (a) Each player in X broadcasts a random bit. Let X_0 denote the set of players who broadcast 0, and X_1 denote the set of players who broadcast 1.
   (b) If $|X_0| \leq |X|/2$, then $X \leftarrow X_0$. Otherwise, $X \leftarrow X_1$.
3. $L \leftarrow X$.

Proposition 6 *When k is an exact power of 2 and $n < 2k$, the simplified protocol elects a good leader with probability at least $p \simeq k^{-(\log k)/4}$.*

Proof: Let $k = 2^t$. With probability roughly $1/\sqrt{k}$, the good players will split evenly between X_0 and X_1. Then, regardless of how the bad players split, the lightest bin will contain a majority of good players, and this majority is an exact power of two. Continuing this argument for t rounds, there is probability roughly $p \simeq \Pi_{i=0}^{t-1}\sqrt{2^i/k} \simeq k^{-(\log k)/4}$ that the lightest bin at round t contains exactly one player, and that this player is good. $\square$

Observe that the above proposition holds when $n = 2k - 1$, and then $\delta = 1/n = 1/(2k-1)$. For this case we have that $p = \delta^{O(\log 1/\delta)}$. We shall show that a similar protocol achieves similar success probability for general values of δ.

The simplified protocol does not work for general values of k. Consider for example a case when S contains three good players and two bad players. If the bad players throw their balls into different bins, then the lightest bin will contain at least one bad player and at most one good player. In the next round, the bad player can broadcast 0 and prevent the good player from being elected.

To overcome the above difficulty, while also generalizing the protocol to arbitrary committee size c, we define the two argument function Half such that $\mathrm{Half}(n, c)$ for $n > c$ is an integer approximately equal to $n/2$, and $\mathrm{Half}(n, c) = c$ modulo $c + 1$. Specifically, if we write n as $2(c + 1)i + j$ with $i \geq 1$ and $-(c + 1) \leq j \leq c$ then $\mathrm{Half}(n) = (c + 1)i - 1$. We now describe the *lightest bin* protocol in full detail.

Lightest Bin (LB) Protocol:

1. $X \leftarrow S$.
2. Repeat while $|X| > c$:
   (a) Each player in X broadcasts a random bit. Let X_0 denote the set of players who broadcast 0, and X_1 denote the set of players who broadcast 1.
   (b) If $|X_0| \leq \mathrm{Half}(|X|, c)$, then $X \leftarrow X_0$. Otherwise, $X \leftarrow X_1$.
3. $L \leftarrow X$.

As a concrete example of how the protocol runs, consider again the case when S contains three good players and two bad players and we wish to choose a leader. Hence $c = 1$ and $\mathrm{Half}(5, 1) = 1$. Assume that in the first round, exactly one of the three good players broadcasts 0. This happens with probability $3/8$. Then if both bad players broadcast 1, the set X_0 continues to the next round, and as it contains just the good player, a good player is chosen as leader. If at least one bad player broadcasts 0, then the set X_1 continues to the next round, even though it may be larger than the set X_0. The set X_1 contains two good players and at most one bad player. Thereafter, if exactly one of the two good players broadcasts 0 (which happens with probability $1/2$), the bad player cannot prevent a good player from being declared as leader. Hence a good leader is chosen with probability at least 3/16.

If in the above example the goal would have been to choose a committee of size 2, then we would have had $\mathrm{Half}(5,2) = 2$. If the good players do not all broadcast the same bit (which happens with probability 3/4), then a good committee is elected already in the first round.

3.1 Analysis of success probability

Lemma 7 *If $k > n/(c+1)$ then the LB protocol selects a good committee with probability $(1/k)^{O(\log k)}$.*

Proof: Observe that a committee of size c is good if more than a fraction of $1/(c+1)$ of the players are good.

We consider for each round the invariant that a fraction of more than $1/(c+1)$ of the players passing to the next round are good. When $\mathrm{Half}(|X|) = (c+1)i - 1$, the invariant is preserved whenever exactly i good players broadcast 0. As the number of good players in the beginning of the round is at least roughly $2i$, and as we may assume that their number is actually not larger than $2i$ (by ignoring some of the good players), then this event happens with probability $\Omega(1/\sqrt{i})$. As there are at most $1 + \log(n/(c+1)) \leq 1 + \log k$ rounds and noting that $i \leq k$, the lemma is proved. $\square$

Lemma 8 *Consider electing a committee of size c when the good players have advantage δ. There is some universal $\delta_0 > 0$ such that for every $0 < \delta < \delta_0$ and for every $n > c/\delta^4$, if the LB protocol is performed only until the number of players is reduced from n to c/δ^4, then with probability at least $1/2$ the fraction of good players remains above $1/(c+1)$.*

Proof: Recall that $k \geq (1+\delta)n/(c+1)$. Let n_i (k_i) be the number of players (good players, respectively) remaining after round i. Then $n_{i+1} \leq n_i/2 + c/2$. When k_i is sufficiently large (we only consider $k = \Omega(1/(\delta_0)^4)$), then with probability at least $1 - q_i \geq 1 - 1/k_i$, $k_{i+1} \geq k_i(1/2 - (k_i)^{-1/3})$ (the good players are partitioned in two by the Binomial distribution, which is centered around its mean). Let t be such that $n_t \leq c/\delta^4$. Then $n_t \leq n2^{-t} + c$, implying $t \simeq \log(k\delta^4)$. Hence if k_i is split in two in each round, then $k_t \simeq 1/\delta^4$, which is large by our choice of small δ_0. Assuming inductively that the imbalance in the splits is always at most $(k_i)^{2/3}$, with probability at least $1 - \sum_{i=0}^{t-1} q_i \geq 1 - \sum_{i=0}^{t-1} 1/k_i > 1/2$, $k_t > k2^{-t}(1 - 2(k_{t-1})^{-1/3})$. As $k_{t-1} > 1/2\delta^4$ it follows that $k_t > n_t\frac{1+\delta}{c+1}(1 - O(\delta^{4/3})) > n_t/(c+1)$, for small enough δ (forced by the choice of δ_0). $\square$

Theorem 9 *Regardless of the number of players, if the advantage is at least δ, then the probability p that the final outcome L of the LB protocol contains at least one good player is at least $\delta^{O(\log 1/\delta)}$.*

Proof: If $\delta \geq \delta_0$ of Lemma 8, then change δ to δ_0 and the proof below then shows that there is a universal success probability $p_0 > 0$, independent of the value of δ. If $\delta < \delta_0$, then Lemma 8 implies that with probability at least $1/2$ the protocol gets to a stage where there are less than c/δ^4 players left, a majority of which are good. Thereafter, the Theorem follows from Lemma 7. □

Recall that $p(n, k, 1)$ is the probability of choosing a good leader when there are k good players. Similarly, let $r(n, k)$ denote the probability of the less likely outcome of a global coin flip when there are k good players. It was known that leader election implies collective coin flipping. Using Theorem 9, we show a stronger connection between the two problems.

Corollary 10 *The collective coin problem and leader election have the same success probability, up to some universal constant. That is, $p(n, k, 1) = \Theta(r(n, k))$.*

Proof: When $k \leq n/2$, neither leader election nor collective coin flipping have robust protocols (e.g., by Theorem 4). Hence we assume $k > n/2$.

Leader election implies collective coin flipping: The elected leader can flip the coin. Hence $r(n, k) \geq p(n, k, 1)/2$.

Collective coin flipping implies leader election: Use the lightest bin protocol to elect a committee of size two. Then use a global coin to select one member of the committee as the leader. Hence $p(n, k, 1) \geq p(n, k, 2)r(n, k)$. For $k > n/2$, we have a committee election problems with parameter $\delta \geq (n/2)(3/n) - 1 \geq 1/2$. Theorem 9 implies that in this case $p(n, k, 2) = \Omega(1)$, implying the corollary. □

The LB protocol is oblivious to the value of δ, which can be an arbitrary positive function of n. Previous studies on leader election focused on the case of fixed $\delta > 0$, and n tending to infinity.

When δ is known to be relatively large compared to $1/n$, then it is possible to condense several rounds of the LB protocol into one round. This leads to a protocol that takes $\log^* n + O(\log 1/\delta)$ rounds, as follows. Each players broadcasts in one round the bits for t consecutive rounds of the LB protocol. This partitions the players into $l = 2^t$ bins, and the players in the smallest bin are chosen to continue the protocol. More generally, we assume that there are l bins, where l need not be a power of two. Each player has its own ball which it throws into a random bin. If l is small enough so that each bin is expected to have roughly the same number of good players, then the analysis of this variant of the protocol is similar to that of Theorem 9. Using this approach, the number of active players can be reduced from n to $O((\log n)^c)$ for some $c > 0$ (e.g., $c = 4$) in a single round, with only negligible loss in the advantage and in the success probability. Iterating this for $\log^* n$ rounds, reduces the number of players to $O(\delta^{-c})$, after which the normal LB protocol is resumed. (Alternatively, the protocol can then be completed in one round, at the expense of worse dependence of p on δ.)

4 The monotone circuit game

As we have seen in Theorem 9, whenever there is a majority of good players, the lightest bin protocol elects a size two committee that with constant probability has at least one good player. We may implement other tasks by presenting two player protocols for them. Specifically, we shall be interested in having the committee of size two choose a leader from the n original players. We show that this two phase approach gives leader election protocols with higher success probability that the LB protocol by itself.

For the leader election task, we may assume that there are two players, one of which is bad, and n candidates from which the two players need to choose a leader. The majority of the candidates are good, and we want to maximize the probability of choosing a good leader.

The protocol we suggest is based on monotone circuits for majority. For simplicity of the presentation, we shall concentrate on circuits with very regular structure. A monotone circuit of depth d is a full binary tree of depth d. The leaves of the tree are labeled by variables and by the constants 0 and 1. Several leaves may have the same label. The internal nodes of the tree are labeled by *and* if they are in an even layer and by *or* if they are in an odd layer. When variables get Boolean values, the circuit computes a monotone function in a natural way, and the output is obtained at the root of the tree. We say that the circuit computes the majority function if the output agrees with the value of the majority of the variables (assume for simplicity that the number of variables is odd). The sorting network of [1] implicitly gives a construction of a majority circuit of depth $O(\log n)$, with a rather large constant hidden by the O notation. Valiant [15] shows that for a circuit of depth roughly $5.3 \log n$, there is a way of labeling its leaves so that it computes the majority function on n variables. His proof is nonconstructive in the sense that it does not describe an explicit labeling of the leaves.

Theorem 11 *If there are circuits for majority of depth d, then there are leader election protocols with success probability at least $p(1/2, 2)2^{-(d+1)/2}$.*

Proof: First choose a committee of size two. As the majority of players are good, Theorem 9 implies that the probability $p(1/2, 2)$ of having at least one good committee member is bounded below by some universal constant. Now call the n candidates x_1 to x_n. Treat them as inputs to a depth d majority circuit. Now the two players play the following game on the circuit. One of the players is the *and* player and the other is the *or* player. The game proceeds in rounds. Starting from the root of the tree, the players trace a path to one of the leaves, by having the players alternate in choosing the next edge on the path. At *and* gates, the *and* player chooses one of the two incoming edges, and at *or* gates the *or* player makes this choice. When a leaf is reached, its label is examined. If the label is a variable, then the respective candidate is selected leader. If the label is 0 (or 1) then the *and* player (*or* player, respectively) is elected leader. (In the more general case where the players themselves are not candidates, then this elected leader can choose a leader at random from the set of candidates.)

We now show that when at least one of the two committee members is good, then the strategy of choosing the next edge at random selects a good candidate with probability at least $2^{-(d+1)/2}$. Assume that the *or* player is good. Treat each of the variables of the good candidates as if it has value 1. Then the output of the majority circuit is 1. When tracing a path from root to leaf, we want to maintain the invariant that the value of the gate at the current location is always 1. If this holds at a leaf, then a good candidate (or the *or* player itself) is selected. Consider now an arbitrary internal node on the path, and assume inductively that it has value 1 (which is true for the root). If it is an *and* node then necessarily the value of the next node is also 1, as both inputs to the node have value 1. If it is an *or* node then maybe only one of its inputs has value 1. But a random choice by the *or* player has probability 1/2 of maintaining the invariant. As there are at most $(d+1)/2$ *or* nodes on the path, the invariant is maintained throughout the execution of the protocol with probability at least $2^{-(d+1)/2}$.

If the *and* player is good, treat each of the variables of the good candidates as if it has value 0, and proceed as above, using duality of 0/1 and of *and/or*. $\square$

Using Valiant's monotone circuits for majority, Theorem 11 implies that $p(n) \geq \Omega(n^{-2.65})$. The monotone circuit approach can be modified so as to obtain the following improvements:

- Higher success probability.
- The protocol can be made explicit.
- The success probability can be expressed as a function of δ.

One needs to observe that Valiant in his proof [15] shows the following amplification result:

Theorem 12 *Let T be a full binary alternating* and/or *tree with* or *gates at the level closest to the leaves, let $\alpha = 1-2(3-\sqrt{5})n/(n-1) \simeq 0.24$, and let δ be sufficiently small in absolute value, in particular satisfying $-1 \leq \delta \leq 1$. Label the leaves independently at random with 0 with probability $\alpha + (1-\alpha)(1-\delta)/2$ and with 1 with probability $(1-\alpha)(1+\delta)/2$. Then* *if the depth of T is $3.3\log(1/\delta) + 2t$ then with probability $1 - 2^{-2^t}$ the circuits outputs 1 if $\delta > 0$ and 0 if $\delta < 0$.*

Theorem 12 has the following implication for two player selection protocols. Given a leader election problem with advantage δ, if the two players could agree on a truly random labeling for the leaves of a circuit of depth $(1 + o(1))3.3\log(1/\delta)$, where a leaf is labeled 0 with probability α and by a random candidate otherwise, this circuit could be used in the proof of Theorem 11. The truly random labeling can be relaxed to agreeing on a somewhat random labeling, provided that the probability of hitting a set of labelings of measure 2^{-2^t} is low (e.g., below one half). Using some encoding mechanism for labelings, the problem of generating a somewhat random labeling can be formulated as a problem of generating a somewhat random binary string of length l, where there is a set of strings of small measure that needs to be avoided. Two player sampling protocols for this problem were studied in [8], and the following simple protocol (which we present for completeness) suffices for our purpose (see proof in [8]).

The protocol proceeds in rounds, with the players switching roles in each round. In a single round, a player uniformly selects an l-dimensional binary vector v_i linearly independent of the vectors used in previous rounds and the other player then selects a random bit σ_i. After l rounds, the string selected is the unique l-bit string whose inner product with every v_i is σ_i.

Corollary 13 *There is an explicit protocol for leader election with success probability $p(\delta) = \Omega(\delta^{1.65})$.*

5 Upper bounds

Theorem 14 *The success probability of collective coin flipping protocols tends to 0 as the fraction of faulty players tends to $1/2$. Quantitatively, for every $\beta > 0$, $r(\delta) = O((1/\delta)^{1-\beta})$.*

Proof: In order to prove Theorem 14 we use the conventions below. It is not hard to see that they may adopted without loss of generality. The players are numbered from 1 to n. The protocol proceeds in steps where in each step a single processor broadcasts a single character from a fixed alphabet. The total number of steps T is fixed in advance, and so is the order in which processors broadcast (e.g., in *round robin* fashion). A *random strategy* specifies for each player a probability distribution over the next character to be broadcast, based on all previous characters that were broadcast. (This is known as a *behavioral strategy*, which in games of full information is the most general kind of strategy.) Good players follow the strategy. Bad players do not necessarily follow the strategy, but they do follow the protocol (broadcast a single character when it is their turn to do so).

The bad players are chosen by an adversary before the protocol begins, and thereafter their messages are chosen by the adversary. A 0-adversary (1-adversary) is one that tries to force the outcome z of the coin to 0 (1, respectively). We assume that the number of good players is k and that the total number of players is $n = 2k - 1$, implying $\delta = 1/n$. (For the sake of negative results, for any value of δ that is the inverse of an odd integer, the most difficult case is when $n = 1/\delta$. The case in which $n = c/\delta$ for some integer $c > 1$ can be simulated by having each of $1/\delta$ players play the role of c players.) We shall fix an arbitrary (supposedly optimal) coin flipping protocol and let $r = \min[Pr[z = 0], Pr[z = 1]]$, taken over the worst adversary.

Consider a set of players $B = \{k + 1, k + 2, \ldots, 2k - 1\}$. In our proof, these players are controlled by a 1-adversary. To define its strategy, consider the collection of sets $\mathcal{T} = \{T_1, \ldots, T_k\}$, where $T_i = \{i\} \bigcup B$. For $1 \leq i \leq k$ and $1 \leq t \leq T$, let p_i^t denote the conditional probability that $z = 1$, where probability is taken conditioned on the first t messages actually broadcast in the protocol, and under the assumption that in future messages, all players in T_i follow the random strategy and the other players are controlled by a 0-adversary. We note that $p_i^0 \geq r$, because $|T_i| = k$. Let v^t be the vector $(p_1^t, \ldots, p_k^t)$, and let $|v^t|$ denote its ρ-norm (i.e., $|v^t|^\rho = \sum_{i=1}^k (p_i^t)^\rho$), where the value of $1 < \rho \leq 2$ will be optimized later in the proof. Let $t + 1$ be a step in which a player from B is to broadcast. The strategy of the 1-adversary is to broadcast a character that maximizes the resulting $|v^{t+1}|$ (breaking ties arbitrarily).

We now give a lower bound for $Pr[z = 1]$ when the 1-adversary controls B and follows the adversarial strategy above. We make the following observations:

1. $p_i^0 \geq r$. Hence $|v^0| \geq r k^{1/\rho}$.
2. After step T, either all $p_i^T = 1$ or all $p_i^T = 0$. $z = 1$ implies $|v^T| = k^{1/\rho}$ whereas $z = 0$ implies $|v^T| = 0$.
3. If a player $i \leq k$ broadcasts at time $t + 1$, then $E[p_i^{t+1}] = p_i^t$ (a martingale property).
4. If $i, j \leq k$, $i \neq j$ and player j broadcasts at step $t + 1$, then $p_i^{t+1} \geq p_i^t$ (because the p_i are defined relative to worst case behavior of player j).
5. If $t + 1$ is a step in which a player in B broadcasts, then $|v^{t+1}| \geq |v^t|$ (as v^t is the weighted average of the possible vectors v^{t+1}, and the adversary's strategy maximizes $|v^{t+1}|$).

The most crucial observation we make is the effect of a message by player $i \leq k$ on $|v^t|$. On every coordinate j other than i, $p_j^{t+1} \geq p_j^t$, making nonnegative contribution towards $|v^{t+1}|$. For coordinate i, we have $E[p_i^{t+1}] = p_i^t$, and w.l.o.g., with nonzero variance (otherwise, ignore this step). Now consider $|v^t|^\rho$. For $\rho > 1$, convexity implies that

$E[|v^{t+1}|^\rho] > |v^t|^\rho$. Moreover, the increase in expectation can be quantified as a function of the variance of p_i^{t+1}. Hence as the protocol progresses, $|v^t|^\rho$ is expected to drift to larger and larger values, making it unlikely to ever reach a value of 0, implying that the 1-adversary almost surely causes $z = 1$.

The above sketch of proof is formalized by modeling $|v^t|$ as a submartingale. See Section A in the appendix. □

Acknowledgements

The author is the incumbent of the Joseph and Celia Reskin Career Development Chair. This manuscript was written while visiting Compaq Systems Research Center, Palo Alto, California. I thank Miki Ben-Or, Ravi Boppana, Moni Naor, Babu Narayanan, Mike Saks and David Zuckerman for directing me to related work, for explaining it to me, and for useful discussions.

References

[1] M. Ajtai, J. Komlos and E. Szemeredi. "An $O(n \log n)$ sorting network". *Proc. 15th ACM Symposium on Theory of Computing*, 1983, 1–9.

[2] M. Ajtai and N. Linial. "The influence of large coalitions". *Combinatorica*, 13 (1993), 129–145.

[3] N. Alon and M. Naor. "Coin-flipping games immune against linear-sized coalitions". *SIAM J. Comput.*, 22 (1993), 403–417.

[4] M. Ben-Or and N. Linial. "Collective coin flipping". In *Advances in Computing Research*, S. Micali, ed., vol. 5: Randomness and Computation, JAI Press, Greenwich, CT, 1989, 91–115.

[5] R. Boppana. Private communication.

[6] R. Boppana and O. Narayanan. "Perfect-information leader election with optimal resilience". *SIAM J. Comput.*, to appear.

[7] J. Cooper and N. Linial. "Fast perfect-information leader-election protocols with linear immunity". *Combinatorica*, 15 (1995), 319–332.

[8] O. Goldreich, S. Goldwasser and N. Linial. "Fault-tolerant computation in the full information model". In *Proc. 32nd Symposium on Foundations of Computer Science*, 1991, 447–457.

[9] J. Kahn, G. Kalai and N. Linial. "The influence of random variables on Boolean functions". In *Proc. 29th Annual Symposium on Foundations of Computer Science*, 1988, 24–26.

[10] S. Karlin and H. Taylor. *A First Course in Stochastic Processes (Second Edition).* Academic Press 1975.

[11] R. Ostrovsky, S. Rajagopalan and U. Vazirani. "Simple and efficient leader election in the full information model". In *Proc. 26th Annual ACM Symposium on the Theory of Computing*, 1994, 234–242.

[12] A. Russell, M. Saks and D. Zuckerman. "Lower bounds for leader election and collective coin-flipping in the perfect information model". In *Proc. 31st Annual ACM Symposium on the Theory of Computing*, 1999, 339–347.

[13] A. Russell and D. Zuckerman. "Perfect information leader election in $\log^* n + O(1)$ rounds". In *Proc. 39th Annual Symposium on Foundations of Computer Science*, 1998, 576–583.

[14] M. Saks. "A robust noncryptographic protocol for collective coin flipping". *SIAM J. Discrete Math.*, 2 (1989), 240–244.

[15] L. Valiant. "Short monotone formulae for the majority function". *J. Algorithms*, 5 (1984), 363–366.

[16] D. Zuckerman. "Randomness-optimal oblivious sampling". *Random Structures and Algorithms*, 11:345–367, 1997.

A Proof of new upper bound for collective coin flipping

In this section we complete the proof of Theorem 14. Recall that we are following the evolution in time t of a k-dimensional vector $v^t = (p_1^t, \ldots, p_k^t)$, where each of its entries is bounded between 0 and 1. We consider its ρ-norm $|v^t|$, where $|v^t|^\rho = \sum_{i=1}^k (p_i^t)^\rho$. We shall choose ρ to be a number slightly larger than 1. In every time step, we either have an adversarial move, which produces an arbitrary new v^{t+1} with $|v^{t+1}| \geq |v^t|$, or a random move, which for some coordinate i satisfies $E[p_i^{t+1}] = p_i^t$, and for every other coordinate j satisfies $p_j^{t+1} \geq p_j^t$. Initially, $p_i^0 \geq r$ for every i, and we wish to upper bound the probability q that at time T $|v^T| = 0$, where T may be arbitrarily large.

A.1 Submartingales

We shall use known results about submartingales (see [10], for example). Let y_t be the message broadcast by a player at step t, and consider the quantity $x_t = |v^t|^\rho$. It satisfies:

1. $E[x_t] \leq \infty$.
2. $E[x_{t+1}|y_1, \dots, y_t] \geq x_t$ (because x^ρ is a convex function when $\rho \geq 1$).
3. x_t is a function of $(y_1, \dots, y_t)$.

Hence x_t is a submartingale.

A.2 Some simplifications

As we are interested in upper bounding the probability that $|v^T| = 0$ and the entries of v^T are nonnegative, we can w.l.o.g. make the following simplifying assumptions:

- Adversarial moves may change v^t but leave $|v^t|$ unchanged.
- Random moves at coordinate i leave the other coordinates unchanged.

We then make the following additional assumption:

- For some fixed ϵ (that may depend of k), random moves have the following effect: $p_i^{t+1} = p_i^t + \epsilon$ with probability $1/2$, and $p_i^{t+1} = p_i^t - \epsilon$ with probability $1/2$.

To justify this last assumption, consider an arbitrary coin flipping protocol P. At each step, a good player would choose its next broadcast according to some probability distribution. Approximate this distribution (which may involve irrational probabilities) by a distribution with rational coefficients. Now every event in the protocol has a probability whose denominator is a product of all denominators of all coefficients. Let ϵ be the inverse of this product.

Now we simulate the behavior of protocol P by using only $\pm\epsilon$ steps. If at step $t+1$ player i makes $p_i^{t+1} = p_i^t + a$ with probability $b/(a+b)$ and $p_i^{t+1} = p_i^t - b$ with probability $a/(a+b)$, then instead take a random walk from p_i^t with $\pm\epsilon$ step size until either $p_i^t + a$ or $p_i^t - b$ is hit. This gives the same distribution. If at step $t+1$ player i can give p_i^{t+1} more than two different values, partition these values into two groups (those above p_i^t and those not above), and take a random walk with $\pm\epsilon$ step size until the expectation of one of these groups is hit. Continue recursively within the group.

A.3 Drift versus variance

Based on the simplifications above, and as we shall be considering only the values $|v^t|^\rho$, we can ignore adversarial moves. For other moves, consider $X_t = x_t - x_{t-1}$. Let $M_t = E[X_t|X_2, \ldots X_{t-1}]$. Assume the random move at step t is at coordinate i, and let $a = p_i^{t-1}$. Then

$$M_t = \frac{(a+\epsilon)^\rho + (a-\epsilon)^\rho}{2} - a^\rho$$

M_t is a decreasing function of a when $1 < \delta < 2$ and $a \geq \epsilon$ (the derivative is negative by concavity of $x^{\rho-1}$), so we can bound M_t from below by assuming $a = 1 - \epsilon$. As we can assume that ϵ is arbitrarily small, we use the the first terms of the Taylor expansion to obtain

$$(a \pm \epsilon)^\rho \simeq a^\rho \pm \rho a^{\rho-1}\epsilon + \frac{\rho(\rho-1)}{2} a^{\rho-2}\epsilon^2$$

with arbitrarily high precision. Substituting $a = 1$ we get $M_t \geq \rho(\rho-1)\epsilon^2/2$.

Let $V_t = E[(X_t - M_t)^2|X_2, \ldots, X_{t-1}]$. Then simple manipulations show that with notation as above

$$V_t = \left(\frac{(a+\epsilon)^\rho - (a-\epsilon)^\rho}{2}\right)^2$$

This expression increases with a and hence can be bounded from above (using the Taylor expansion) by $V_t \leq \rho^2\epsilon^2$.

Let $\alpha = (\rho-1)/2\rho$, which is a positive constant whenever $\rho > 1$. ¿From the above we obtain for every step of our protocol $M_t \geq \alpha V_t$.

A.4 An inequality for partial sums

For X_i and α as above we use the following lemma (equation (4.14) in [10]):

Lemma 15 *The probability that the sum of the X_i's ever drops below $-l$ is at most $1/(1+\alpha l)$.*

In our special case, where the values of X_i are bounded (and in fact, arbitrarily small by a small enough choice of ϵ) we have the following corollary:

Corollary 16 *For integer $c > 1$, the probability that the sum of the X_i's ever drops below $-cl$ is at most $1/(1+\alpha l)^c$.*

Proof: When the sum drops below $-l$, we may assume that it is in fact $-l$, because individual changes to the sum are arbitrarily small. Hence to drop below $-cl$, we need c successive drops of magnitude l, and the probability of each new drop is upper bounded independently of previous drops. $\Box$

We can now complete the proof of Theorem 14. Assume that $r > (1/k)^{1-\beta}$ for some $\beta > 0$. Then for some $1 < \rho < 1/(1-\beta)$ we have $x_0 = |v^0|^\rho \geq k(1/k)^{(1-\beta)\rho} > (\log k)/\alpha$, where $\alpha = (\rho - 1)/2\rho$ as above. Observe that $x_T = 0$ only if $\sum X_i$ drops below $-(\log k)/\alpha$, which has probability at most $2^{-\log k} = 1/k$, by Corollary 16. Hence we exhibited a strategy for the adversary that causes the coin to come up 1 with probability $1 - 1/k$, contradicting the assumed value of r.

B Some specialized upper bounds

B.1 One round coin flipping

One round protocols are of special interest. There we assume that the good players broadcast simultaneously, and then the bad players broadcast their messages. We make no restrictions on the length of a message. Let $r(2k-1,k)$ denote the probability of the less likely outcome of a global coin flip z when there are k good players and $k-1$ bad players. The protocol in which each player sends a random bit and the value of z is the majority of the bits has $r(2k-1,k) = 2^{-k}$. This is best possible for one round protocols.

Theorem 17 *For every one round protocol* $r(2k-1,k) \leq 2^{-k}$.

Proof: For set S with $|S| = k$ and player $i \in S$, call a message m by i *deadly* for S if broadcasting m leaves S in a situation where regardless of the messages broadcast by the other members of S, the players outside of S can force $z = 0$. Let $q(i,S)$ denote the probability that i broadcasts a deadly message for S.

We now distinguish between two cases.

Case 1: The expectation over i, S satisfies, $E[q(i,S)] \geq 1/2$. In this case there is some set S with $|S| = k$ such that the expectation over its players i, $E[q(i,S)] \geq 1/2$. Then (by comparing geometric and arithmetic mean) with probability at least $1 - 2^{-k}$ some player in S broadcasts a deadly message for S. When this happens, the complement of S (which is of size $k-1$) can force $z = 0$, and hence $Pr[z = 1] \leq 2^{-k}$.

Case 2: The expectation over i, S satisfies, $E[q(i,S)] \leq 1/2$. In this case there is some set S with $|S| = k-1$ such that the expectation over the k players i outside of S, $E[q(i, S \bigcup \{i\})] \leq 1/2$. Then with probability at least $1 - 2^{-k}$ some player i outside S broadcasts a message that is not deadly for the respective $S \bigcup \{i\}$. When this happens, S (which is of size less than $k-1$) can force $z = 1$. Hence $Pr[z = 0] \leq 2^{-k}$. $\Box$

For one round leader election protocols, see Section C.3.

An interesting open question regarding collective coin flipping protocols is whether there are one round protocols with success probability lower bounded as a function of the advantage δ, independent of the number of players n.

B.2 Two player selection games

We used two player games as a subroutine for leader election protocols.

Proposition 18 *For every two player protocol of selecting one out of $2k-1$ candidates of which k candidates are good, one of the players has a strategy by which the probability of choosing a good candidate is at most $1/k$.*

Proof: We assume for simplicity that the protocol is sure to end. Player a has a strategy of forcing the selected candidate to be between 1 and k, as these may be the good candidates. When player b plays randomly against this strategy, one candidate $1 \leq i \leq k$ has probability at most $1/k$ of being chosen. Now if b is the good player and $\{i, k+1, k+2, \ldots, 2k-1\}$ are the good candidates, then player a has a strategy underwhich the probability of choosing a good candidate is at most $1/k$. $\Box$

C Some toy examples

In this section we present selection protocols for several toy examples. This may help in avoiding making unfounded conjectures regarding what cannot be done.

C.1 Leader election with one faulty player

The following protocol elects a good leader with probability $1-1/n$ when there is just one faulty player. Hence Theorem 2 is not applicable in this setting.

Start *Baton Passing* at player 1. The penultimate player to receive the Baton then chooses player 1 as the leader with probability $1/n$, and the player never receiving the baton as the leader with probability $1-1/n$.

C.2 Leader election with five players

The following leader election protocol chooses a good leader with probability at least $4/9$ when the number of good players is three and the number of bad players is two. The question of determining the best value of $p(5,3,1)$ is open.

Player 1 removes a player chosen uniformly at random. The removed player then removes one player: player 1 with probability $1/2$, any other player with probability $1/6$. Observe that regardless of whether player 1 is good or bad, with probability $2/3$ one of the two removed players is bad. On the remaining three players, select a leader using a protocol that succeeds with probability $2/3$ when two players are honest, as discussed above. The overall success probability is $(2/3)^2 = 4/9$.

Note that for the corresponding selection problem (five players, at least three of which are good, need to select one of five candidates, at least three of which are good), there are always two players who can cause a bad candidate to be chosen with probability $(2/5)^{1-2/5} \simeq 0.577 > 5/9$ (by Theorem 2).

C.3 One round leader election

Finding the best one round protocol for leader election is open even in the simplest case of two good players and one bad one. The lightest bin protocol chooses a good leader with probability $1/2$ (when the good players go into different bins), but is not optimal. Let $p = (\sqrt{5}-1)/2$, $q = (3-\sqrt{5})/2$, so that $p+q=1$ and $q = p^2$. The following protocol elects a good leader with probability $p \simeq 0.618$.

Players A and B each send 0 with probability q, and 1 with probability p. Player C does not speak. If player A sent 0, then player B is leader. If players A and B sent 1, then player C is leader. If player A sent 1 and B sent 0, then player A is leader.

If all players are honest, then A, B and C each have probability pq, q and $p^2 = q$ of being leader, respectively. Neither B nor C can increase their own probability of being leaders by cheating. If player A cheats and deterministically sends 1, then A's probability of being leader increases to q.

Remark: Recently, together with Ehud Wieder it was shown that the above protocol is indeed optimal for the one round three player case. For the four player case (three good and one bad) Ryan O'Donnel shows a one round protocol for which a good leader is chosen with probability $3/4$.