Triadic Consensus: A Randomized Algorithm for Voting in a Crowd

ASHISH GOEL and DAVID LEE, Stanford University

Motivated by the hope of developing group consensus mechanisms over the internet, we study an urn-based voting rule where each participant acts as a voter and a candidate. We prove that when participants lie in a one-dimensional space, this voting protocol finds a $(1 - \epsilon/\sqrt{n})$ approximation of the Condorcet winner with high probability while only requiring an expected $O(\frac{1}{\epsilon^2} \log^2 \frac{n}{\epsilon^2})$ comparisons on average per voter. Moreover, this voting protocol is shown to have a quasi-truthful Nash equilibrium: namely, a Nash equilibrium exists which is not truthful, but produces a winner with the same probability distribution as that of the truthful strategy. We discuss implications of our work for designing low communication complexity and 'truthful' voting protocols.

Categories and Subject Descriptors: F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms: Algorithms, Economics, Theory

Additional Key Words and Phrases: Computational social choice, probabilistic voting rules, crowdsourcing, quasi-truthful voting, strategic voting, communication complexity

ACM Reference Format:
Goel, A., Lee, D. T. 2012. Triadic Consensus: A randomized algorithm for voting in a crowd. ACM X, X, Article X (February 2012), 19 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

Voting is often used as a method for achieving consensus among a group of individuals. This may happen, for example, when a committee chooses a representative or some friends go out to watch a movie. When the group is small, this process is relatively easy; however, for larger groups, the typical requirement of ranking all candidates becomes impractical and rough heuristics are often applied to narrow down opinions to a few representative ones before a vote is taken.

This problem of large-scale preference aggregation is even more interesting in light of the rising potential of crowdsourcing. For instance, suppose a university were to pose the following question to all its alumni, students and faculty: "What is an ideal curriculum for computer science undergraduates?" Or suppose a city government asked its constituencies: "What is the ideal budget that also cuts 50 percent of the deficit?" Indeed, the second question is not too far from what has been discussed in the wake of the recent federal budget stalemates. Several websites, such as Widescope[1], have sprung up around the idea that crowdsourcing may even be useful as a tool towards solving political questions.

However, while it is easy to solicit proposals to these questions, aggregation seems challenging. In particular, it may not be practical for a participant to even look through the resulting set of proposals, making seemingly simple tasks such as choosing top ranked proposals, difficult.

In this paper, we propose a randomized voting rule for the previously mentioned consensus scenarios. In our problem setting, each participant proposes exactly one proposal, representing his or her stance on the question of interest. A random triad of participants is then selected and each selected member is made to vote between the other two. Roughly speaking (details are elaborated in Section 2), if there is a three-way tie, the participants are thrown out from the election; otherwise, the losers are

[1]widescope.stanford.edu

replaced by 'copies' of the winner. This process is then repeated until there is only a single participant remaining, who is declared the winner.

We show that for single dimensional participant spaces, Triadic Consensus converges approximately to the Condorcet winner[2] with high probability, while only requiring an average of $\sim \log^2 n$ (conjectured to be $\sim \log n$) comparisons per individual. As we will point out in Section 4.2.1, this approximation is quite reasonable for this application. For instance, suppose participants are random samples from $[0, 1]$. Then if the true median were used to select the winner, it would be distributed with approximately the same probability distribution as that of Triadic Consensus.

In addition, in the same one-dimensional setting, Triadic Consensus is *quasi-truthful* for concave utilities. In particular, even though we show in Section 2.1 that the truthful strategy *is not* a Nash equilibrium, a different Nash equilibrium exists which chooses the winner *with the same probability distribution* as in truthful voting. Surprisingly, we achieve this result by counterintuitively allowing voters to express cyclical preferences (e.g. $a > b$, $b > c$, and $c > a$). Though this may seem to allow additional strategies for manipulations, it is shown to also allow strategies that prevent manipulation.

1.1. Related Work and Our Contributions

Given the long history of work on voting theory, it is not surprising that the problems we tackle have been, for the most part, thought about before. Here, we give a brief overview of related work, followed by a summary of our contributions (Section 1.1.5). For in-depth reading, we refer the reader to Brandt et al. [2012].

1.1.1. Voting rule criteria. One of the earliest criterion introduced for evaluating voting rules is known as the Condorcet criterion, introduced by Marquis de Condorcet[3]. It states that if a candidate exists who would win against every other candidate in a majority election, then this candidate should be elected. Unfortunately, such a candidate does not always exist. Since then, many other criteria have been introduced as ways to evaluate voting rules. However, in the surprising result known as Arrow's Impossibility Theorem, Arrow [1950] proved that there were three desirable criterion that no deterministic voting rule could satisfy. This was expanded by Pattanaik and Peleg [1986] to show that a similar result holds for probabilistic voting rules.

1.1.2. Strategic manipulation. This sparked a wave of impossibility results, including the classical Gibbard-Satterthwaite Impossibility Theorem. Define a voting rule to be *strategy-proof* if it is always in a voter's interest to submit his true preference, regardless of the other voter rankings. Gibbard [1973] and Satterthwaite [1975] independently showed that all deterministic voting rules that were strategy-proof must either be dictatorships or never allow certain candidates to win. This was extended to show that among all probabilistic voting rules, only very simple voting rules were strategy-proof[Gibbard 1977].

Numerous attempts at circumventing these impossibility result have been made. Bartholdi et al. [1989] first proposed using computational hardness as a barrier against manipulation in elections. However, despite many NP-hardness results on manipulation of voting rules[Faliszewski and Procaccia 2010], it was shown that there do not exist any voting rules that are *usually* hard to manipulate[Conitzer and Sandholm 2006].

[2]The Condorcet winner is a candidate who would beat any other candidate in a pairwise majority election. In single dimensional spaces, this happens to be the median participant.
[3]See Young [1988] for a fascinating historical description of the early work of Condorcet.

Procaccia [2010] used the simple probabilistic voting rules of Gibbard [1977] to approximate common voting rules in a strategy-proof way, but the approximations are weak and they show that, for many of these voting rules, no strategy-proof approximations can be much stronger. Birrell and Pass [2011] extended this idea to approximately strategy-proof voting, proving that there exist tight approximations of any voting rule that are close to strategy-proof. Recently, Alon et al. [2011] studied a special case of approval voting and showed that even though no deterministic strategy-proof mechanism has a finite approximation ratio, a randomized strategy-proof mechanism exists which has a good approximation ratio.

1.1.3. Communication complexity. When the number of candidates is large, it is also important to study voting rules from the perspective of the burden on voters. This problem was studied by Conitzer and Sandholm [2005] in the context of communication complexity. In their work, they study the worst case number of bits that voters need to communicate in order to determine the ranking or corresponding winner of common voting rules; for many of these voting rules, it was shown that the number of bits required is essentially the same as what is required for reporting the entire ranking. In addition, Conitzer and Sandholm [2002] showed another disappointing result: for many common voting rules, determining how to elicit preferences efficiently is NP-complete, even in the case when perfect knowledge about voter preferences is assumed. These previous results are all for exact rank or winner determination. Lu and Boutilier [2011] proposed several algorithms for elicitation under approximate winner determination. Though no theoretical guarantees are given, they support their algorithms with experimental simulations.

1.1.4. Single-peaked preferences. One special case that avoids the many discouraging results above is that of single-peaked preferences[Black 1948] (or other domain restrictions). Single-peaked preferences are those for which candidates can be described as lying on a line. Every voter's utility function is peaked at one candidate and drops off on either side. For such preferences, a Condorcet winner always exists and is the candidate who is the median of all voter peaks. This winner can be found by the classical median voting rule, which has each voter provide their peak and returns the median peak. It turns out that the median voting rule is both strategy-proof[Moulin 1980] and has a low communication complexity of $O(n \log m)$[Escoffier et al. 2008], where n is the number of voters and m is the number of candidates. Conitzer [2009] also studies the problem of eliciting voter preferences or the aggregate ranking using comparison queries.

It is useful to point out that the strategy-proof result relies primarily on restricting the allowed input from the voter, i.e. restricting the domain of the voting rule. Similarly, it should be noted that the low communication complexity of the median voting rule relies on the ability to calculate the median of peaks, which comes from an assumed knowledge of the underlying axis. When this axis is unknown, Escoffier et al. [2008] provide an $O(mn)$ algorithm that finds and returns the axis.

1.1.5. Our contributions. We introduce the following novel concepts:

(1) *A localized consensus mechanism for large groups*. We propose Triadic Consensus as an approach for large groups to make decisions using small decentralized decisions among groups of three, where the only central processing required is in random number generation (alternatively, random mixing in a group of participants).
(2) *Quasi-truthful voting rules and cyclical preferences*. When each participant is a voter and a candidate, we demonstrate that allowing participants to express cyclical preferences ($a > b$, $b > c$, and $c > a$) can introduce strategies that detect and

protect against strategic manipulation (as shown in Section 4). We call a voting rule that admits such a strategy quasi-truthful since they achieve truthful results with votes that are not, strictly speaking, truthful.

In addition, we are able to apply these ideas to overcome some previous limitations of rules for single-peaked preferences. Specifically, Triadic Consensus *does not assume prior knowledge of an axis* and *does not restrict voter preference domains*, but still has the following properties:

(1) It finds a $(1-\epsilon/\sqrt{n})$ approximation of the Condorcet winner with high probability[4].
(2) It has a communication complexity of $O(\frac{n}{\epsilon^2}\log^2\frac{n}{\epsilon^2})$ for the above approximation factor, i.e. $\sim n\log^2 n$ (conjectured to be $\sim n\log n$) for a $1-\frac{1}{\sqrt{n}}$ approximation.
(3) It has a quasi-truthful Nash equilibrium when participants have concave utility functions.

The fact that these proofs do not rely on the previously mentioned assumptions are important since they provide directions for finding generalizations to higher dimensional spaces. Indeed, there may be many situations in which proposals may lie in high dimensional spaces, but are approximately single dimensional (small second singular value); in such a situation, it would not be possible to assume an axis.

As a contrast to our results, consider the classical median voting rule which is strategy-proof, gives an exact Condorcet winner, and has $O(n\log m)$ time. However, since it assumes prior knowledge of an axis and restricts voter preference domains, there are no possible generalizations to higher dimensional spaces. Also, when prior knowledge of an axis is unknown, the calculation of this axis will increase the time complexity to $O(mn)$ time. When we consider the median voting rule from a practical perspective, it is not clear how to ask people to “state their peak”. The most typical example of single-peaked preferences is in two-party politics, which has a conservative to liberal axis. But what does it mean for a participant to say, “I am 70 percent republican”?

Triadic Consensus solves this practical question by using pairwise comparisons: participants simply state their opinions on policies and then make pairwise comparisons on which opinions they prefer.

1.1.6. Outline of the paper. Before continuing, we briefly describe the structure of the remaining sections. In Section 2, we detail Triadic Consensus and introduce the notion of quasi-truthfulness. This is followed by Section 3, in which we give a technical overview of the main intuitions and techniques of our proofs. The main body of the approximation and communication complexity proofs will be contained in Section 4, which will be followed by the main body of the quasi-truthfulness proofs in Section 5. Finally, we will conclude with a discussion on future directions in Section 6.

2. TRIADIC CONSENSUS AND QUASI-TRUTHFULNESS

Triadic Consensus is applicable in a setting with n individuals, each of which has a candidate solution to the problem of interest and an internal ranking over all the proposed candidates. Formally, each individual is represented as a point x in some space X and his or her preference ranking is determined by a distance metric $d(x, y)$ on X. If $d(x, y_1) \leq d(x, y_2)$, then x will rank y_1 higher than y_2; that is, x prefers proposals that are closer to him. Our proofs will apply to the case when X is the set of reals $\mathbb{R}$

[4]A precise statement of this can be found in Section 4.

ALGORITHM 1: Triadic Consensus

Input: An urn with k labeled balls for each participant $1, 2, \ldots, n$
Output: A winning candidate i.
if $nk \equiv 0 \bmod 3$ **then**
  Duplicate one ball in the urn uniformly at random;
while *there is more than one label* **do**
  Sample three balls (with labels x, y, z) uniformly at random with replacement;
  $w = \text{TriadicVote}(x, y, z)$;
  if $w = \emptyset$ **then**
    Remove the three sampled balls from the urn;
  else
    Relabel all the sampled balls with the winning label w;
return Remaining label;

ALGORITHM 2: TriadicVote

Input: Candidates x, y, z
Output: One of $\{x, y, z\}$ if there is a winner, $\emptyset$ otherwise
if *two of more of x, y, z have the same id* **then**
  return the majority candidate;
x votes between y and z;
y votes between x and z;
z votes between x and y;
if *each received exactly one vote* **then**
  return $\emptyset$;
else
  return the candidate with two votes;

and $d(x, y)$ is the Euclidean distance, but the algorithm is well defined for the general setting[5].

The best way to understand Triadic Consensus (Algorithm 1) is to imagine an urn with balls, each of which is labeled by a participant id. The urn starts with k balls for each of the n participants[6]. Then, at each step, the algorithm samples three balls uniformly at random (with replacement) and performs a TriadicVote (Algorithm 2) on the three corresponding participants.

If the three participants x, y, and z are unique, the TriadicVote subroutine consists of a single comparison for each of the selected participants: x votes between y and z, y between x and z, and z between x and y. These votes can be distributed in some permutation of $2, 1, 0$ or split $1, 1, 1$. In the first case, the participant who received two votes is returned as the winner. In the second case, a tie (represented as $\emptyset$) is returned. If two or more of the selected ids are the same, i.e. are the same person, then he or she is automatically returned as the winner without any votes[7].

If a winner was returned from the TriadicVote, then the three balls are relabeled with the winning id and placed back into the urn; otherwise, the three balls are removed. This process is repeated until there is only one participant id remaining, which

[5]Slight modifications are needed when there can be truthful ties.
[6]Increasing k makes the approximation tighter, but requires more comparisons to converge, so choosing k is essentially choosing a tradeoff between approximation and time.
[7]This can be viewed as the result of a vote where each of the majority id votes for the other which means one of them must win.

ALGORITHM 3: Truthful voting

Input: Voter x, candidates y, z
Output: One of $\{y, z\}$
if $d(x, y) \leq d(x, z)$ **then**
  return y;
else
  return z;

ALGORITHM 4: Quasi-truthful voting

Input: Voter x, candidates y, z
Output: One of $\{y, z\}$
if $d(y, z) \geq d(y, x)$ *and* $d(y, z) \geq d(z, x)$ **then**
  return y or z according to some strategy (such as the one in Section 3.3);
else
  return a truthful comparison between y and z;

is declared the winner. It is helpful to point out that the action of removing balls can be thought of as a deterrent for manipulation. This is because, in the case of candidates in a Euclidean space, a three way tie should rarely occur if participants are voting truthfully[8].

There is one remaining caveat to the algorithm. Notice that the total number of balls in the urn remains constant except for the case of a tie, in which the total number of of balls is reduced by three. Then, if the initial number of balls kn is a multiple of three, there is some possibility of ending up with no balls in the urn before a winner is declared. To avoid this, we randomly choose a participant at the beginning of the algorithm and start him off with $k + 1$ labeled balls.

2.1. Truthfulness and Quasi-truthfulness

Consider a TriadicVote between participants x, y, and z. Recall that truthful voting means that each participant will vote for the other closest participant. If participants do vote truthfully (Algorithm 3), then it is fairly easy to understand what will happen; a voter x will simply vote for whichever candidate out of y and z lies closer to his or her own position. Unfortunately, voting truthfully is not a Nash equilibrium as illustrated in the following example.

Example 2.1. Four participants lie at positions 0, 5, 6, and 7. Suppose participants 0, 5, and 7 are selected for a TriadicVote. Since they are voting truthfully, 0 votes for 5, 5 votes for 7, and 7 votes for 5. As a result, 5 wins and the resulting urn consists of three balls for 5 and one token for 6.

Now suppose participant 7 did not vote truthfully and voted for participant 0. Then there would be a three-way tie, which means that all selected balls are eliminated, leaving an urn with only participant 6. Clearly, participant 7 would prefer this second non-truthful situation.

However, Triadic Consensus admits another notion of truthfulness. Imagine a TriadicVote between x, y, and z, and assume without loss of generality, that x would have been the winner if everyone voted truthfully. This can only occur if both y and z vote

[8]The only possible exception is when candidates are located on the points of an equilateral triangle. Though we do not face this issue in 1D, this can be addressed by allowing participants to express a tie in the TriadicVote.

for x. In such a situation, it does not matter who x votes for; since x has already received two votes, he will win regardless of whether he votes for y or z. We call such voting behavior quasi-truthful (Algorithm 4) since the result is the same as truthful voting. We will show that when participants have utility functions that are concave in the distance of a given proposal, x can use his vote to disincentivize y and z from manipulating the TriadicVote. For the time being, we will illustrate it for the same example.

Example 2.2. Four participants lie at positions 0, 5, 6, and 7, all of whom would like to minimize the expected distance of the winning proposal to their position. Now suppose participants $0, 5$, and 7 are selected for a TriadicVote. A quasi-truthful strategy would be for 0 to vote for 5, 5 to vote for 0 and 7 to vote for 5. Since 5 got two votes, he wins and the resulting urn consists of three balls for 5 and one token for 6.

Now suppose participant 7 deviates from this strategy and votes for 0. Then participant 0 gets two votes and he wins. The resulting urn consists of three balls for 0 and one token for 6, which is clearly worse for participant 7. Likewise, suppose participant 0 deviates from this strategy and votes for 7. Then there is a three-way tie and all selected balls get eliminated. The resulting urn consists of a single token for 6, which is clearly worse for participant 0.

In this example, even though participant 5 did not vote truthfully, the end result of the TriadicVote was the same as if truthful votes were cast. In addition, no players can benefit from deviating from this strategy.

3. TECHNICAL OVERVIEW AND INTUITION

In this section, we provide a proof sketch for the main results in Sections 4 and 5, and also describe intuitions for these results.

3.1. Approximation of the Condorcet Winner

The primary idea in proving the approximation result is to reduce the Triadic Consensus urn to previously known results for fixed size urns. Suppose we color all balls with participant ids 1 to i red and all balls with participant ids $i+1$ to n blue. Then each time we perform a TriadicVote, there is some probability of recoloring a red ball blue or a blue ball red which depends on a function of the current fraction of red and blue balls. For this particular urn with $k = 1$, we can apply known theorems[Lee and Bruck 2012] to derive the probability $\mathbf{Pr}[w \leq i]$ of the winner being one of the red balls 1 to i. Then the probability that participant i is the winner is simply $\mathbf{Pr}[w \leq i] - \mathbf{Pr}[w \leq i-1]$, which is surprisingly the binomial distribution $\mathrm{Bin}(n-1, \frac{1}{2})$, i.e.

$$\mathbf{Pr}[w = i] = \left(\frac{1}{2}\right)^{n-1} \binom{n-1}{i-1}$$

Following this, some algebraic computations show that uniformly duplicating a ball does not change this expression. For the general case, a similar argument would yield the binomial distribution $\mathrm{Bin}\left(nk-1, \frac{1}{2}\right)$ for the probability that a given ball wins[9].

Note that the standard deviation of $\mathbf{Pr}[w = i]$ is $\frac{1}{2}\sqrt{n}$. This means that if there are 1000 participants, then Triadic Consensus will produce a winner between $[468, 532]$ with ninety-five percent probability. As we explain more rigorously in Section 4.2.1, we should not work too hard to reduce this magnitude of error. If these 1000 participants are random samples from an underlying distribution, the exact Condorcet winner would be distributed with the same magnitude of error away from the true Con-

[9]This is not technically well-defined. However, as we point out in Section 4, we can use this expression for the winning probability of a *ball* without changing the overall winning probability of participants.

dorcet winner of the underlying distribution. We also note in Section 4.2.2 that Triadic Consensus has the property of quickly eliminating outliers since each participant must convince two others to vote for him. This is stated more formally by a simple comparison to another randomized voting rule that only uses pairwise comparisons without the use of triads.

3.2. Communication Complexity

To derive the communication complexity results, we use the same reduction to fixed size urns. Let $\mathbb{E}[T]$ be the expected time before all balls are either on participants 1 to i or on participants $i+1$ to n. By applying known theorems[Lee and Bruck 2012] on the expected time of convergence for the corresponding urn with $k=1$, we get the bound

$$\mathbb{E}[T] \leq n \ln n + O(n)$$

Note that this bound is true for all i. Then we can use it as a bound for the expected time to halve unique participants ids. After $\log n$ of these iterations, there can only be one remaining participant who must be the winner, which means that Triadic Consensus takes $O(n \log^2 n)$ comparisons when $k = 1$. Simulations indicate that the true complexity is only $O(n \log n)$ comparisons. The bound we provide is loose because each time the participant ids are halved, we are allowing for the case when remaining balls are still uniformly distributed among the remaining participants. Again, these results can be trivially extended to the general case of k balls per participant by using nk instead of n in the expression for $\mathbb{E}[T]$. We also note that even though the average number of votes per participant is $O(\log n)$, participants that 'survive' longer may cast significantly more votes.

In Appendix A.1, we will show some simulations for two-dimensional spaces indicating that Triadic Consensus still achieves tight approximations with $O(n \log n)$ communication complexity when participants are located in a plane.

3.3. Quasi-truthful Nash Equilibrium

The intuitions for finding a quasi-truthful Nash can be understood by generalizing Example 2.2. Suppose that three participants $x < y < z$ are drawn for a TriadicVote and x votes for y, z votes for y, and y votes for x. In this case, y wins and participants x and z are the only participants who can change the result by a strategic vote.

— If x votes strategically for z, then there will be a three-way tie instead of a win for y.
— If z votes strategically for x, then x will win instead of y.

It seems intuitively bad for z if x wins over y since x is strictly farther away. If this is true, then the only participant who has any hope to manipulate the election is x, the person that y voted for. From this intuition, we get the strategy expressed in Algorithm 5.

First, we note that Algorithm 5 relies on the existence of a participant y or z who would prefer a win for x rather than a three-way tie (when x should win in truthful voting). This is where the requirement of concave utilities comes in. It turns out that when all players have a utility function that is concave in the distance of a proposal to his position, then at least one of the two losing players x and z prefers a win for y over a three-way tie. Then, we are able to prove that Algorithm 5 is a Nash equilibrium with the following proof by induction.

Base Case: Algorithm 5 is a Nash equilibrium for $n = 1, 2$ *balls.* This is trivial since no strategic moves can be played when there is only one or two balls.

ALGORITHM 5: Quasi-truthful Nash for concave utilities

Input: Voter x, candidates y, z
Output: One of $\{y, z\}$
if $d(y,z) \geq d(y,x)$ *and* $d(y,z) \geq d(z,x)$ **then**
  if y *would prefer a win for* x *rather than a three-way tie in a truthful world* **then**
    return y;
  else if z *would prefer a win for* x *rather than a three-way tie in a truthful world* **then**
    return z;
else
  return a truthful comparison between y and z;

Inductive Step: Assume that Algorithm 5 is a Nash equilibrium for $n-3$ balls. Then it is also a Nash equilibrium for n balls. We will prove this statement in two smaller steps:

(1) Assume Algorithm 5 is a Nash equilibrium for $n-3$ balls. Then for any TriadicVote with participants $x < y < z$ in an urn with n balls, if y votes for x in the the Nash equilibrium (Algorithm 5), then x maximizes his expected utility by voting for y rather than z, i.e. x prefers a win for y rather than a three-way tie.
(2) Assume Algorithm 5 is a Nash equilibrium for $n-3$ participants and that the above statement is true. Then we show that Algorithm 5 is a Nash equilibrium for n participants. This is done by defining a comparison relation between urns that formalizes the intuition we had that x prefers a win for y over a win for z. With this definition, we can define a coupling of two urns: one in which x plays the optimal strategy, and one in which x plays according to Algorithm 5. We show that for every coupled history, the urn from Algorithm 5 does at least as well as the optimal urn in expected utility. This means that Algorithm 5 is also an optimal strategy for x, which proves that it is a Nash equilibrium for urns with n balls.

Then by carrying out the Inductive Hypothesis from $1, 2$ to $4, 5$ to $7, 8$, etc..., we get the result for all $n \not\equiv 0 \bmod 3$. For $n \equiv 0 \bmod 3$, we note that a random ball is duplicated at the beginning, which puts it into the case of $n \equiv 1 \bmod 3$.

ALGORITHM 6: Simple Quasi-truthful Nash with modified Triadic Consensus

Input: Voter x, candidates y, z
Output: One of $\{y, z\}$
if $d(y,z) \geq d(y,x)$ *and* $d(y,z) \geq d(z,x)$ **then**
  if *It is the first TriadicVote* **then**
    return His preferred candidate;
  else
    return His less preferred candidate;
else
  return a truthful comparison between y and z;

At first glance, the Quasi-truthful Nash in Algorithm 5 seems impractical since a voter must know the utilities of all other participants. There is a simple fix to Triadic Consensus that solves this problem. When there is a three way tie, instead of throwing away all three balls, Triadic Consensus will ask for another vote. If there is another three way tie, then we will throw the three balls out; otherwise, we will replace the balls with the winner. Now there is a simple Quasi-truthful Nash: the candidate that should win will simply vote truthfully during the first TriadicVote. If there is a cycle and another TriadicVote is taken, he should then switch his vote.

The proofs for this mechanism follow directly from the proofs for the original mechanism, so we will primarily refer to the original in the following proof sections. The main bridge between the two proofs is to note that if there is at least one voter who prefers the truthful winner over a three-way tie, then when he is voted for, he will ensure that there is not a three-way tie by voting truthfully. The other voter cannot create a cycle, so he should also vote truthfully.

4. TRIADIC CONSENSUS APPROXIMATES THE CONDORCET WINNER WITH LOW COMMUNICATION COMPLEXITY IN 1D

In this section, we will detail the proof sketches from Section 3.

4.1. Background: Switching probability and convergence time of fixed size urns

We first briefly describe results on fixed size urns that we will need. A fixed size urn contains some number of balls, which are each colored either red or blue. Let R_t and B_t be the number of red and blue balls respectively at time t, where $R_t + B_t = n$. Also, let $p_t = \frac{R_t}{n}$ denote the fraction of red balls. At every discrete time t, either a red ball is sampled with probability $f(p_t)$, a blue ball is sampled with probability $f(1-p_t)$, or nothing happens with the remaining probability. The function $f : [0,1] \to [0,1]$ is called an urn function and satisfies $0 \leq f(x) + f(1-x) \leq 1$ for $0 \leq x \leq 1$. If a ball was sampled, it is then recolored to the opposite color and placed back into the urn. This process repeats until some time T when all the balls are the same color, i.e. $R_T = n$ or $R_T = 0$.

We will later show that Triadic Consensus is closely related to fixed size urns with urn function $f(p) = 3p(1-p)^2$. For now, we simply state the following theorems that come from general ones stated in Lee and Bruck [2012][10].

THEOREM 4.1. *Let a fixed size urn start with R_0 red balls out of n total balls and have an urn function $f(p) = 3p(1-p)^2$. Let T denote the first time when either $R_T = n$ or $R_T = 0$. Then,*

$$\mathbf{Pr}[R_T = n] = \left(\frac{1}{2}\right)^{n-1} \sum_{j=1}^{R_0} \binom{n-1}{j-1}$$

THEOREM 4.2. *Let a fixed size urn start with R_0 red balls out of n total balls and have an urn function $f(p) = 3p(1-p)^2$. Let T denote the first time when either $R_T = n$ or $R_T = 0$. Then,*

$$\mathbb{E}[T] \leq n \ln n + O(n)$$

4.2. Approximation of the Condorcet Winner

When the space of candidates X is the reals $\mathbb{R}$, we show that any quasi-truthful strategy will result in a close approximation of the Condorcet winner, which is the median candidate.

LEMMA 4.3. *Let x, y, and z be three unique participants which lie on a line. Then the winner of TriadicVote(x, y, z) will be the median participant.*

PROOF. WLOG, suppose $x < y < z$. Then x and z vote for y who is the median participant. Since y gets two votes, she wins. □

[10]Theorem 4.2 requires some algebra that may not be immediately clear from the general theorem. For the convenience of the reader, we include these calculations in Appendix C.

THEOREM 4.4. *Let there be n participants in $\mathbb{R}$. Label them as participants $1, 2, \ldots n$ from $-\infty$ to ∞ (leftmost position to rightmost position) and let w denote the winning id after running Triadic Consensus with $k = 1$. Then assuming that participants vote quasi-truthfully*[11]*,*

$$\mathbf{Pr}[w = i] = \left(\frac{1}{2}\right)^{n-1} \binom{n-1}{i-1}$$

PROOF.

Case 1: $n \not\equiv 0 \bmod 3$. Since n is not a multiple of three, Triadic Consensus will start with exactly one ball per participant. Also, since all participants are voting quasi-truthfully, the TriadicVote will never return a tie. Our proof will be by a reduction to the fixed size urn described in Theorem 4.1.

Color balls with ids $1, 2, \ldots, i$ red and balls with ids $i+1, i+2, \ldots n$ blue. Each time three balls are sampled, we can have the following four color combinations: Three red; two red and one blue; one red and two blue; and three blue. When all are the same color, then the relabeling scheme will not change the colors of the balls. When two are red and one blue, we can apply Lemma 4.3 to claim that the winner must be one of the red balls. Likewise, when one is red and two are blue, the winner must be one of the blue balls. Letting p_r and p_b denote the fraction of red and blue balls respectively, we note the following four possible results:

Three red	With probability p_r^3, there is no change in colors.
Two red, one blue	With probability $3p_r^2 p_b$, one blue ball is replaced by a red ball.
One red, two blue	With probability $3p_r p_b^2$, one red ball is replaced by a blue ball.
Three blue	With probability p_b^3, there is no change in colors.

This is exactly the dynamics described by a constant population urn with i red balls, $n-i$ blue balls, and urn function $f(p) = 3p(1-p)^2$. Now, note that $w \le i$ iff the winning ball is red. Then by applying Theorem 4.1, we have

$$\mathbf{Pr}[w \le i] = \left(\frac{1}{2}\right)^{n-1} \sum_{j=1}^{i} \binom{n-1}{j-1}$$

To finish the proof, we note the the probability that the winner is participant i is simply the probability that the winner is one of $1, 2, \ldots, i$ and not one of $1, 2, \ldots, i-1$:

$$\mathbf{Pr}[w = i] = \mathbf{Pr}[w \le i] - \mathbf{Pr}[w \le i-1] = \left(\frac{1}{2}\right)^{n-1} \binom{n-1}{i-1}$$

Case 2: $n \equiv 0 \bmod 3$. In this case, we will uniformly select one participant d to duplicate. To calculate the new value of $\mathbf{Pr}[w = i]$ for a given i, we notice that d can fall into one of three cases: $d < i, d = i, d > i$. In each of these cases, we get an instance of the problem with $n+1$ candidates. If $d < i$, then participant i wins if the $i+1$-th ball wins. If $d = i$, then participant i wins if either the i-th or $i+1$-th ball wins. Finally, if $d > i$, then participant i wins if the i-th ball wins. Then letting w_n denote the winner

[11]This theorem can also be applied to the general case of n balls. Defining the probability that the i-th ball wins with the given probability will give the correct expression for the winning probability of participants.

of Triadic Consensus with n participants, we have that for $n \equiv 0 \bmod 3$,

$$\begin{aligned}\mathbf{Pr}[w_n = i] &= \frac{i-1}{n}\mathbf{Pr}[w_{n+1} = i+1] + \frac{1}{n}\mathbf{Pr}[w_{n+1} = i \text{ or } i+1] + \frac{n-i}{n}\mathbf{Pr}[w_{n+1} = i] \\ &= \frac{i}{n}\mathbf{Pr}[w_{n+1} = i+1] + \frac{n-i+1}{n}\mathbf{Pr}[w_{n+1} = i] = \left(\frac{1}{2}\right)^{n-1}\binom{n-1}{i-1}\end{aligned}$$

which finishes the proof. □

Using standard probabilistic arguments[Motwani and Raghavan 1995], we get the following corollary.

COROLLARY 4.5. *Let there be n participants in $\mathbb{R}$ and let w denote the winning id after running Triadic Consensus with $k = O(\frac{1}{\epsilon^2}\log\frac{1}{\delta})$. Then assuming that participants vote quasi-truthfully, w will be a $(1-\epsilon/\sqrt{n})$ approximation of the Condorcet winner with probability at least $1-\delta$.*

4.2.1. How bad is $O(\sqrt{n})$ error?. We would like to point out that for crowdsourcing applications, participants can be viewed as noisy samples from some underlying distribution. Then, even using a single ball per candidate ($k = 1$) may not be that bad. The noise in the participant sampling process may already result in similar errors.

For example, suppose that the participants are drawn independently and uniformly from $[0, 1]$ so that the true Condorcet candidate would be one with position $\frac{1}{2}$. Suppose $\frac{1}{2}$ lies between the k-th and $k+1$-th participant and that we had an oracle that told us this index k. Then k is clearly binomially distributed, which is the same distribution as the case of Triadic Consensus with $k = 1$. In other words, the true Condorcet winner has a standard deviation of $\frac{1}{2}\sqrt{n}$ participants between him and the sampled Condorcet winner.

4.2.2. Triadic Consensus eliminates outliers quickly. Triadic Consensus has the intuition of quickly eliminating outliers since each participant needs to convince two other participants to vote for him in order to win. We will give a brief comparison to a more standard form of random comparisons. Consider Hot-or-Not Consensus, in which two balls are randomly chosen as candidates and one single ball is randomly chosen as the voter. The voter then votes between the two chosen balls and the two candidate balls are replaced with the winning candidate. For a continuous distribution of participants, we have the following one-step analysis.

THEOREM 4.6. *Let a continuous distribution of voters be uniformly distributed between zero and one. Let $g_{\textit{Hot-or-Not}}(x)$ and $g_{\textit{Triadic}}(x)$ be the probability density of x being the next winning candidate in Hot-or-Not and Triadic Consensus respectively. Then, $g_{\textit{Triadic}} = 6x(1-x)$ and $g_{\textit{Hot-or-Not}} = 3x(1-x) + \frac{1}{2}$. In particular,*

$$g_{\textit{Hot-or-Not}} = \frac{1}{2}g_{\textit{Triadic}} + \frac{1}{2}g_{\textit{Unif}}$$

where $g_{\textit{Unif}}$ is the uniform distribution over the interval $[0, 1]$.

PROOF. The proof for this can be found in Appendix A.2. □

In other words, a more naive algorithm (Hot-or-Not Consensus) can be thought of as a mix between Triadic Consensus and the (really bad) method of randomly picking a candidate[12].

[12]This is only an intuition based on a one-step comparison and should not be interpreted as a comparison of their final approximation values.

4.3. Communication Complexity

When the space of candidates X is the reals $\mathbb{R}$, we show that any quasi-truthful response to Triadic Consensus will have, on average, a sublinear communication complexity.

THEOREM 4.7. *Assuming quasi-truthful voting, Triadic Consensus has a total communication complexity of* $O(kn\log^2(kn))$.

PROOF. Equation 4.2 is an upper bound on the time it takes to halve the number of remaining participants. Initially, we color all balls with participant ids $1, 2, \ldots, \frac{n}{2}$ red and all balls with participant ids $\frac{n}{2}+1, \frac{n}{2}+2, \ldots, n$ blue. Then it takes an expected time of $n \ln n + O(n)$ for all balls to be either red or blue. At this point, all the remaining balls must have participant ids that are either between 1 and $\frac{n}{2}$ or $\frac{n}{2}+1$ and n. Suppose WLOG, that the remaining participant ids lie between 1 and $\frac{n}{2}$. Then we can recolor the balls for participants 1 to $\frac{n}{4}$ red and the balls for participants $\frac{n}{4}+1$ to $\frac{n}{2}$ blue. Again, it takes another $n \ln n + O(n)$ expected time for all the balls to converge to one of these groups. After $\log n$ of these iterations, there can be at most one remaining participant, so we are done. Therefore, the total number of TriadicVotes is bounded by

$$(\log n)\mathbb{E}[T] \approx \log n(kn \log kn)$$

We get our final result by noting that there can be at most three comparisons per TriadicVote to get our final result. □

As previously mentioned, the above theorem is a very coarse analysis of the communication complexity. In reality, at the first time when all balls are either in $\left[1, \frac{n}{2}\right]$ or $\left[\frac{n}{2}+1, n\right]$, the balls are not uniformly distributed across the interval of $\frac{n}{2}$ participants. Rather, the balls are very tightly packed around the center, giving the following conjecture:

CONJECTURE 4.8. *Assuming quasi-truthful voting, Triadic Consensus has a total communication complexity of* $O(kn \log kn)$.

We are also hopeful that Triadic Consensus has a low communication complexity for high dimensional spaces. In Appendex A.1, we show some simulations in 2D that also demonstrate low communication complexity.

5. TRIADIC CONSENSUS HAS A QUASI-TRUTHFUL NASH EQUILIBRIUM IN 1D

When the space of candidates X is the reals and participants have a concave utility function, we demonstrate that Algorithm 5 is a quasi-truthful Nash equilibrium. The skeleton of the proof is described in Section 3. Here, we flesh out the details for proving the two parts in the described Inductive Step. Because of space constraints, we will simply give proof sketches for the more algebraic proofs, leaving the details to Appendix B.

5.1. Assume that Algorithm 5 is a Nash equilibrium for $n-3$ balls and $x < y < z$ WLOG. Then for x as the only deviator from Algorithm 5, if y votes for x, then x does not increase his expected utility by voting for z instead of y

Note that a vote from x for z would result in a three-way tie, whereas a vote from x for y would result in a win for y. Since Algorithm 5 states that y votes for the player who prefers a win for y over a three way tie, this statement is already satisfied *assuming such a player always exists for* y *to vote for*. Then we just need to show that for an urn with n balls, at least one of x and z will have a higher utility in the case of a win for y as compared to the case of a three-way tie. The proof can be described in two parts:

(1) Part One: We show that we only need to consider urns where all the balls lie in the interval $[x, z]$, i.e. a proof for this specific situation would extend to all urn configurations.
(2) Part Two: We show that if x and z have concave utility functions in the distance of a proposal, then in urns where all balls lie in the interval $[x, z]$, at least one of x and z prefers a win for y over a three-way tie.

5.1.1. Preliminaries. Let R denote the urn resulting from a three-way tie. Then urn R must have $n-3$ balls since we removed one ball from each of x, y, and z. By our assumption, this means that Algorithm 5 is optimal for the remaining decisions in R.

Let S denote the urn resulting from a win for y. To get urn S from R, we just need to add three balls to participant y. Our approach will be to show that the quasi-truthful strategy in urn S will give a higher expected utility for at least one of x and z. If this is true, then we are done since it is clear that this participant would want y to win.

We index the balls in urn S as $b_1, b_2, \ldots, b_n$ from the leftmost participant position to the rightmost and let b_l, b_{l+1}, b_{l+2} denote the three balls on participant y that do not exist in urn R. For each ball b_i and quasi-truthful voting, there is some probability that the ball wins in urn R (if balls b_l, b_{l+1}, b_{l+2} had not been added) and some probability that the ball wins in urn S (after balls b_l, b_{l+1}, b_{l+2} have been added)[13]. We would like to see how this probability changes. Denote this difference in the winning probability of b_i as $\Delta p(b_i) = \mathbf{Pr}[b_i \text{ wins in urn } S] - \mathbf{Pr}[b_i \text{ wins in urn } R]$.

We use $U_x(b_i)$ and $U_z(b_i)$ to denote the utilities of x and z respectively gained from a win for ball b_i. Define $U_x(b_i)$ to be a concave utility function if $U_x(b_i) = f_x(d(b_i, x))$ where f_x is a concave function and $d(b_i, x)$ is the distance between b_i and x. Note that f_x must necessarily be monotonically decreasing since candidates prefer proposals that are closer to them. We will use U_x^S and U_z^S to denote the expected utility for x and z respectively from quasi-truthful voting in an urn S. Finally, we use $\Delta U_x = U_x^S - U_x^R$ and $\Delta U_z = U_z^S - U_z^R$ to denote the difference in expected utility in quasi-truthful voting for urn S as compared to urn R.

Then our goal is now reduced to showing that one of $\Delta U_x \geq 0$ and $\Delta U_z \geq 0$ is true.

5.1.2. Part One: We only need to consider urns where all balls lie in the interval $[x, z]$. We first show that the difference in winning probability $\Delta p(b_i)$ is positive in an interval around the added balls b_l, b_{l+1}, b_{l+2} and negative everywhere else[14]. Specifically,

LEMMA 5.1. *Suppose Triadic Consensus is run on urns R and S as defined in Section 5.1.1. Then for quasi-truthful voting,*

$$\Delta p(b_i) > 0 \textit{ if } \min(l, n/2) \leq i \leq \max(l+2, n/2)$$
$$\Delta p(b_i) < 0 \textit{ otherwise}$$

PROOF. The proof consists primarily of algebraic manipulations of Theorem 4.4. For the sake of space constraints, the proof is included in Appendix B. □

With this, we can now prove the main lemma for Part One.

LEMMA 5.2. *Let R and S be urns as defined in Section 5.1.1. From R, create a new urn R' by moving all balls left of x to x and all balls right of z to z. Similarly, from S, create a new urn S' by moving all balls left of x to x and all balls right of z to z. Let $\Delta U'_x = U_x^{S'} - U_x^{R'}$ and $\Delta U'_z = U_z^{S'} - U_z^{R'}$. Then,*

$$\Delta U_x \geq \Delta U'_x \quad \textit{and} \quad \Delta U_z \geq \Delta U'_z$$

[13]Recall that we can use Theorem 4.4 to describe the probability that a *ball* wins without changing the winning probabilities of participants.
[14]Since balls b_l, b_{l+1}, b_{l+2} do not exist in urn R, we just use probability 0 for $\mathbf{Pr}[b_i \text{ wins in urn } R]$.

PROOF. For the sake of space constraints, we will only sketch the proof. A full detailed version can be found in Appendix B. The basic idea is that we will break this proof into three cases.

Case 1: $b_{n/2} \in [x, z]$ for urn S. When the median ball is located in $[x, z]$, we know from Lemma 5.1 that $\Delta p(b_i) < 0$ for all balls that are not in the interval $[x, z]$. For any balls with $\Delta p(b_i) < 0$, ΔU_x and ΔU_z are decreased when the balls are moved closer to x and z respectively so long as the ball index is unchanged. But since moving the balls left of x to x brings the balls closer to both x and z, we must have decreased both ΔU_x and ΔU_z. The same argument holds for moving balls right of z to z.

Case 2: $b_{n/2} < x$ for urn S. When the median ball is located left of x, we know from Lemma 5.1 that $\Delta p(b_i) < 0$ for all balls right of z and left of $b_{n/2}$. For the same reason as in Case 1, when we move these balls closer to x and z, we can only decrease ΔU_x and ΔU_z. This allows us to move all balls right of z to z and all balls left of $b_{n/2}$ to $b_{n/2}$. Now, for the balls right of $b_{n/2}$ and left of x, we have that $\Delta p(b_i) > 0$. For such balls, we decrease ΔU_x and ΔU_z by moving them farther away from x and z. Putting these together, we are now able to move all balls left of x to $b_{n/2}$.

Now the second trick is to note that all balls left of x are now at the same position (that of $b_{n/2}$). This allows us to treat them as a single ball b^* with $\Delta p(b^*) = \sum_{i=1}^{k-1} \Delta p(b_i)$, where k is the index of the first ball right of x. We then show that $\Delta p(b^*) < 0$, which allows us to apply the same arguments as above to move all these balls to position x and decrease ΔU_x and ΔU_z.

Case 3: $b_{n/2} > z$ for urn S. The proof is symmetric to that of Case 2. □

At this point, we have shown that $\Delta U_x \geq \Delta U'_x$ and $\Delta U_z \geq \Delta U'_z$. Then if we are able to claim that one of $\Delta U'_x$ and $\Delta U'_z$ is greater than zero, it is obviously true that one of ΔU_x and ΔU_z is greater than zero. This finishes Part One since we have shown that we only need to prove our statement for urns where all balls lie in $[x, z]$.

5.1.3. Part Two: Suppose x and z have concave utility functions. Then for urns R and S with all balls in $[x, z]$, one of $U_x \geq 0$ and $U_z \geq 0$ holds

THEOREM 5.3. *Let R and S be urns as defined in Section 5.1.1 and let all their balls lie within the interval $[x, z]$. Then if participants x and z have concave utility functions, at least one of $\Delta U_x \geq 0$ and $\Delta U_z \geq 0$ is true.*

PROOF. The detailed version of this proof can be found in Appendix B. Since it is primarily algebraic and also notation-heavy, we will only demonstrate it here with an example.

Let S be an urn with five balls: b_1 is located at position x; b_2, b_3, and b_4 are located at position y; and b_5 is located at position z. Then urn R is an urn with the two balls b_1 and b_5. Let participants x and z have utility functions $U_x(b_i) = f_x(d(b_i, x))$ and $U_z(b_i) = f_z(d(b_i, z))$ respectively, where f_x and f_z are monotonically decreasing and concave functions.

We know (see Theorem 4.4) that for balls b_1, b_2, b_3, b_4, and b_5, $\mathbf{Pr}[b_i \text{ wins in urn } S]$ is $\frac{1}{16}, \frac{4}{16}, \frac{6}{16}, \frac{4}{16}$, and $\frac{1}{16}$ respectively; $\mathbf{Pr}[b_i \text{ wins in urn } R]$ is $\frac{1}{2}, 0, 0, 0$, and $\frac{1}{2}$, respectively;

which means that $\Delta p(b_i)$ is $-\frac{7}{16}, \frac{4}{16}, \frac{6}{16}, \frac{4}{16}$, and $-\frac{7}{16}$, respectively. Then,

$$\begin{aligned}\Delta U_x &= \Delta p(b_1)U_x(b_1) + \left[\sum_{i=2}^{4} \Delta p(b_i)U_x(b_i)\right] + \Delta p(b_5)U_x(b_5)\\ &= -\frac{7}{16}f_x(d(x,x)) + \frac{14}{16}f_x(d(x,y)) - \frac{7}{16}f_x(d(x,z))\\ &= -\frac{7}{16}[f_x(d(x,x)) - f_x(d(x,y))] - \frac{7}{16}[f_x(d(x,z)) - f_x(d(x,y))]\end{aligned}$$

Since $\frac{7}{16}[f_x(d(x,x)) - f_x(d(x,y))] \geq 0$, we have that

$$\Delta U_x \geq 0 \iff \frac{-\frac{7}{16}[f_x(d(x,z)) - f_x(d(x,y))]}{-\frac{7}{16}[f_x(d(x,x)) - f_x(d(x,y))]} \geq 1$$

Using similar arguments and the fact that $\frac{7}{16}[f_z(d(z,x)) - f_z(d(z,y))] \leq 0$, we have that

$$\Delta U_z \geq 0 \iff \frac{-\frac{7}{16}[f_z(d(z,z)) - f_z(d(z,y))]}{-\frac{7}{16}[f_z(d(z,x)) - f_z(d(z,y))]} \leq 1$$

By using concavity and Lemma C.3 in Appendix C,

$$\begin{aligned}\frac{f_x(d(x,z)) - f_x(d(x,y))}{f_x(d(x,x)) - f_x(d(x,y))} &\geq \frac{(z-x)-(y-x)}{(x-x)-(y-x)} = \frac{z-y}{x-y}\\ &= \frac{(z-z)-(z-y)}{(z-x)-(z-y)} \geq \frac{f_z(d(z,z)) - f_z(d(z,y))}{f_z(d(z,x)) - f_z(d(z,y))}\end{aligned}$$

This means that if $\frac{z-y}{x-y} \geq 1$, then $\Delta U_x \geq 0$. Otherwise, if $\frac{z-y}{x-y} \leq 1$, then $\Delta U_z \geq 0$. Obviously, one of these must be true, so we are done. □

5.2. Assuming that Algorithm 5 is a Nash equilibrium for $n-3$ balls and given the result of Section 5.1, then Algorithm 5 must be a Nash equilibrium for n balls

Suppose that x is the only possible deviator from Algorithm 5. We want to show that with the inductive hypothesis given for $n-3$ and the results of Section 5.1, then x cannot increase his expected utility by deviating from Algorithm 5.

Our proof strategy will be to use a coupling argument. Let OPT denote the optimal strategy for x. We consider two urns R and S. In urn R, x plays according to Algorithm 5.1. In urn S, x plays according to OPT. Now, we will couple the TriadicVote's of these urns in the following way:

(1) Let $r_1, r_2, \ldots, r_n$ denote the balls in urn R as indexed from leftmost position to rightmost position. Let $s_1, s_2, \ldots, s_n$ denote the balls in urn S as indexed from leftmost position to rightmost position.
(2) Then for every TriadicVote, when balls r_i, r_j, r_k are randomly drawn from urn R, balls s_i, s_j, s_k will be drawn from urn S.

To compare the evolution of these two urns, we define the following notion.

Definition 5.4. Given two urns R and S, each with n balls, number the balls in R from left to right as $r_1, r_2, \ldots, r_n$ and number the balls in S from left to right as $l_1, l_2, \ldots, l_n$. Then R x-dominates S if

$$\begin{aligned}s_i &\leq r_i \text{ for } i : r_i < x\\ s_i &= r_i \text{ for } i : r_i = x\\ s_i &\geq r_i \text{ for } i : r_i > x\end{aligned}$$

In words, if R x-dominates S, this means that for any ball left of x in R, the corresponding ball in S must be farther left of x. Likewise, for any ball right of x in R, the corresponding ball in S must be farther right of x. Finally, for all balls that are exactly at x in R, the corresponding ball in S must also be exactly at x.

We note that initially, urns R and S are identical, which also satisfies the definition of R x-dominates S. Then we show that after a single TriadicVote, the resulting urns R' and S' will still satisfy the property of R' x-dominates S'.

THEOREM 5.5. *Let urns R and S be coupled together as defined above. Then letting R' and S' be the resulting urns after the TriadicVote in R and S respectively, we must have R' x-dominates S'.*

PROOF. We note that since R x-dominates S, whenever x participates in a TriadicVote in R, it is also participating in a TriadicVote in S. Moreover, if x is the left, middle, or right participant in the TriadicVote of urn R, then it is also the left, middle, or right participant respectively in the TriadicVote of urn S. Then we have four possible cases:

Case 1: x is not one of r_i, r_j, r_k. Since we have assumed that all other participants are voting according to Algorithm 5, r_j wins in urn R and s_j wins in urn S. Note that all other balls have not moved and we also have $d(x, r_j) \leq d(x, s_j)$ and $r_j, s_j > x$. Then R' x-dominates S'.

Case 2: x is two or more of r_i, r_j, r_k. Then x wins in both urns R and S and all the selected balls r_i, r_j, r_k and s_i, s_j, s_k are moved to x. Since all other balls have not moved, R' x-dominates S'.

Case 3: x is the middle participant r_j. Since we have assumed that all other participants are voting according to Algorithm 5, r_i and r_k vote for x in urn R and s_i and s_k vote for x in urn S. Therefore, x wins in both coupled runs and it is clear that R' x-dominates S' since all other balls have not moved.

Case 4: x is one of the side participants (r_i WLOG). Since we have assumed that all other participants are voting according to Algorithm 5, r_k votes for r_j in urn R and s_k votes for s_j in urn S. Also, since x is following Algorithm 5 in R, then r_i votes for r_j, which means that r_j wins in urn R. In urn S, there could be two different outcomes depending on how s_j votes.

(1) Suppose s_j votes for s_k. Then one of s_j, s_k will win in urn S depending on who s_i votes for. But both of these are farther from x than r_j and on the same side of x, i.e. $d(x, r_j) \leq \min(d(x, s_j), d(x, s_k))$ and $r_j, s_j, s_k > x$. Therefore, R' x-dominates S'.
(2) Suppose s_j votes for s_i. Then since x is playing the optimal strategy in urn S, s_i must vote for s_j (by Section 5.1 and the inductive hypothesis for $n-3$). But this means that s_j wins in urn S. Once again, s_j is farther from x than r_j and on the same side of x, i.e. $d(x, r_j) \leq d(x, s_j)$ and $r_j, s_j > x$. Therefore, R' x-dominates S'.

□

We have shown that after one coupled TriadicVote, R' x-dominates S'. However, for Theorem 5.5, we only used the fact that R x-dominates S. Therefore, we can apply the theorem to state that $R^{(t)}$ x-dominates $S^{(t)} \implies R^{(t+1)}$ x-dominates $S^{(t+1)}$, where $R^{(t)}$ and $S^{(t)}$ are the resulting urns from R and S respectively after t TriadicVote's. Then this must be true even after the urns have converged to a winner and we have that the winner of R x-dominates the winner of S. This is equivalent to saying that the winner of R is at least as close to x as the winner of S.

Then, this means that x gains at most the same utility from OPT as compared to the utility gained from Algorithm 5 and we are done.

6. FUTURE DIRECTIONS

There are many future directions and open problems that this work presents:

Triadic Consensus. For the algorithm itself, the primary problem that begs to be worked on is an analysis for higher dimensional or even non-Euclidean spaces. It is an open question whether Triadic Consensus achieves low communication complexity for general preference profiles. The authors believe this is the case; however, even if this is true, it would need to be accompanied with some way to evaluate the winner produced. For example, does it approximate the Condorcet winner whenever it exists? Does Triadic Consensus find central nodes in graphs? Does it approximate the notion of a generalized median in $\mathbb{R}^d$?

Similarly, it would also be exciting to extend the work on quasi-truthfulness to higher dimensional spaces. The authors do not believe that a naive extension will suffice; however, it seems possible that probabilistic strategies coupled with other punishments for manipulation will be able to achieve this goal.

Truthful voting rules. When participants are voters and candidates, we have indicated that manipulation can often be detected. It would be interesting to use this idea, possibly along with the theme of triads, quasi-truthfulness, and cyclic preferences, to design truthful voting rules. For example, one could imagine the following variant of the Borda count: for each of the $\binom{n}{3}$ triads, add one point to the score of the winner[15].

Communication complexity. Another exciting problem is to make new approximate and randomized voting rules that have low communication complexity. In particular, it would be useful to have a voting rule where the maximum number of comparisons per voter is small (say, $O(\log n)$). In Triadic Consensus, only the average number of comparisons is small, which may still prevent it from being widely applicable to large internet crowdsourcing applications. Indeed, if such a voting rule could be designed, it may be another approach to tackle strategic voting in the sense that participants are only comparing a vanishingly small number of participants; this may prevent them from being able to determine how to manipulate.

Consensus mechanisms. On the direction of group consensus mechanisms, one possible extension of this work is to bring it outside of voting. Namely, rather than having the randomly selected triads *vote*, it would be interesting to analyze other sorts of dynamics that are more collaborative or game-theoretic.

Urn voting rules. It would be also interesting to study generalized urn voting rules. This could include different ball replacement schemes or even more elaborate generalizations. For example, balls could be labeled with participant *and* proposal ids so that only proposal ids are changed after a TriadicVote. Such urn voting rules are interesting because they can be interpreted as local decisions made by small groups of people and also provide a natural framework for studying (non-trivial) probabilistic voting rules.

[15]The Borda count is equivalent to giving the winner two points, the next highest scoring participant one point, and the loser zero points.

ACKNOWLEDGMENTS

The authors would like to thank Pranav Dandekar for helpful discussions on Triadic Consensus and Vincent Conitzer for pointing us to several pieces of related literature and for posing the idea of using approximations to lower communication complexity.

REFERENCES

ALON, N., FISCHER, F., PROCACCIA, A., AND TENNENHOLTZ, M. 2011. Sum of us: Strategyproof selection from the selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*. 101–110.

ARROW, K. 1950. A difficulty in the concept of social welfare. *Journal of Political Economy 58,* 4, 328–346.

BARTHOLDI, J., TOVEY, C., AND TRICK, M. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare 6*, 227–241.

BIRRELL, E. AND PASS, R. 2011. Approximately strategy-proof voting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*.

BLACK, D. 1948. On the rationale of group decision-making. *Journal of Political Economy 56,* 1, 23–34.

BRANDT, F., CONITZER, V., AND ENDRISS, U. 2012. Computational social choice. In *Multiagent Systems*. MIT Press.

CONITZER, V. 2009. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research 35*, 161–191.

CONITZER, V. AND SANDHOLM, T. 2002. Vote elicitation: Complexity and strategy-proofness. In *Proceedings of the 17th AAAI Conference*.

CONITZER, V. AND SANDHOLM, T. 2005. Communication complexity of common voting rules. In *Proceedings of the 6th ACM Conference on Electronic Commerce (ACM-EC)*. 78–87.

CONITZER, V. AND SANDHOLM, T. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st AAAI Conference*. 627–634.

ESCOFFIER, B., LANG, J., AND ÖZTÜRK, M. 2008. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008)*.

FALISZEWSKI, P. AND PROCACCIA, A. 2010. Ai's war on manipulation: Are we winning? *AI Magazine 31,* 4, 53–64.

GIBBARD, A. 1973. Manipulation of voting schemes: A general result. *Econometrica 41,* 4, 587–601.

GIBBARD, A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica 45,* 3, 665–681.

LEE, D. AND BRUCK, J. 2012. Generalized switching in urns. *To be posted on arXiv*.

LU, T. AND BOUTILIER, C. 2011. Robust approximation and incremental elicitation in voting protocols. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*.

MOTWANI, R. AND RAGHAVAN, P. 1995. *Randomized Algorithms*. Cambridge University Press.

MOULIN, H. 1980. On strategy-proofness and single peakedness. *Public Choice 35,* 4, 437–455.

PATTANAIK, P. AND PELEG, B. 1986. Distribution of power under stochastic social choice rules. *Econometrica 54,* 4, 909–921.

PROCACCIA, A. 2010. Can approximation circumvent gibbard-satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. 836–841.

SATTERTHWAITE, M. 1975. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory 10*, 187–217.

YOUNG, H. P. 1988. Condorcet's theory of voting. *The American Political Science Review 82,* 4, 1231–1244.

Online Appendix to:
Triadic Consensus: A Randomized Algorithm for Voting in a Crowd

ASHISH GOEL and DAVID LEE, Stanford University

A. APPROXIMATION OF THE CONDORCET WINNER AND COMMUNICATION COMPLEXITY

A.1. Simulations in 2D

We show two examples of Triadic Consensus in 2D. In the first case (a grid), we see similar tightness and communication complexity to our results from 1D. In the second case, we try to design a difficult scenario by densely populating the perimeter of a circle and adding a single point at its center. Surprisingly, Triadic Consensus is still able to select this center point with non-trivial probability.

Example A.1. $n = m^2$ voters are placed on the points $(0,0),(0,1),\ldots,(m-1,m-1)$ to form a $m \times m$ grid. The Condorcet winner in this scenario is at the median point $(\frac{m}{2},\frac{m}{2})$. From the simulation results below, we can see that Triadic Consensus picks winners that are closely distributed around the winner. The average number of votes each voter casts is $\sim O(\log n)$.

The data in the table below represent 100 iterations of Triadic Consensus. Column 1 is the dimensions of the grid, column 2 is the average winner, column 3 is the standard deviation of the winner (square root of the mean squared distance from the average winner), column 4 is the average number of votes per participant, and column 5 is the standard deviation of the average number of votes per voter (over the 100 iterations).

	mean winner	σ of winner	mean votes/voter	σ of votes/voter
5x5	(1.96, 1.97)	.953	1.966	0.425
10x10	(4.45, 4.5)	1.157	3.405	0.517
20x20	(9.61, 9.47)	1.236	4.618	0.417
40x40	(19.64, 19.53)	1.594	6.056	0.368
80x80	(39.57, 39.69)	1.555	7.293	0.324

Example A.2. n voters are placed uniformly around a circle (in the plane) with radius 1 and centered at $(0,0)$. A single voter is placed at the point $(0,0)$. The Condorcet winner in this scenario is the point $(0,0)$. Surprisingly, we find that even as the number of points increases on the perimeter, the probability of the randomized algorithm selecting the single Condorcet winner still remains non-trivial. The average number of votes each voter casts is $\sim O(\log n)$.

The data in the table below represent 1000 iterations of Triadic Consensus. Column 1 is the number of participants in the circle, column 2 is the fraction of times that $(0,0)$ is selected as the winner, column 3 is the average number of votes per participant, and column 4 is the standard deviation of the average number of votes per voter (over the 1000 iterations).

	% times (0,0) wins	mean votes/voter	σ of votes/voter
25	0.368	2.225	0.513
100	0.338	4.189	0.775
400	0.305	6.628	1.185
1600	0.295	9.224	1.696
6400	0.302	11.764	2.321

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

A.2. Triadic Consensus eliminates outliers quickly

THEOREM A.3. *Let a continuous distribution of voters be uniformly distributed between zero and one. Let $g_{Hot\text{-}or\text{-}Not}(x)$ and $g_{Triadic}(x)$ be the probability density of x being the next winning candidate in Hot-or-Not and Triadic Consensus respectively. Then, $g_{Triadic} = 6x(1-x)$ and $g_{Hot\text{-}or\text{-}Not} = 3x(1-x) + \frac{1}{2}$. In particular,*

$$g_{Hot\text{-}or\text{-}Not} = \frac{1}{2} g_{Triadic} + \frac{1}{2} g_{Unif}$$

where g_{Unif} is the uniform distribution over the interval $[0, 1]$.

PROOF. For uniformly distributed participants, we have the density function $f(x) = 1$ and cumulative density function $F(x) = x$. In Triadic Consensus, x wins if he is selected along with a candidate to the left and right of him. Then, we have

$$g_{\text{Triadic}}(x) = 3! f(x) F(x)(1 - F(x)) = 6x(1-x)$$

In Hot-or-Not Consensus, x wins against y only if the voter z is closer to x than y. Then,

$$\begin{aligned} g_{\text{Hot-or-Not}}(x) &= 2f(x) \int_0^x f(y) \left(1 - F\left(\frac{x+y}{2}\right)\right) \mathrm{d}y + 2f(x) \int_x^1 f(y) F\left(\frac{x+y}{2}\right) \mathrm{d}y \\ &= 2\left[\int_0^x \left(1 - \frac{x+y}{2}\right) \mathrm{d}x + \int_x^1 \left(\frac{x+y}{2}\right) \mathrm{d}x\right] = 3x(1-x) + \frac{1}{2} \end{aligned}$$

□

B. QUASI-TRUTHFUL NASH EQUILIBRIUM

LEMMA B.1. *Suppose Triadic Consensus is run on urns R and S as defined in Section 5.1.1. Then for quasi-truthful voting,*

$$\begin{aligned} &\Delta p(b_i) > 0 \textit{ if } \min(l, n/2) \le i \le \max(l+2, n/2) \\ &\Delta p(b_i) < 0 \textit{ otherwise} \end{aligned}$$

PROOF. Recall that $\Delta p(b_i) = \mathbf{Pr}[b_i \text{ wins in urn } S] - \mathbf{Pr}[b_i \text{ wins in urn } R]$. From Theorem 4.4, we have that a quasi-truthful strategy in urns R and S give,

$$\mathbf{Pr}[b_i \text{ wins in urn } S] = \left(\frac{1}{2}\right)^{n-1} \binom{n-1}{i-1}$$

$$\mathbf{Pr}[b_i \text{ wins in urn } R] = \begin{cases} \left(\frac{1}{2}\right)^{n-4} \binom{n-4}{i-1} & \text{if } i \le l-1 \\ \left(\frac{1}{2}\right)^{n-4} \binom{n-4}{i-4} & \text{if } i \ge l+3 \\ 0 & \text{if } i = l, l+1, l+2 \end{cases}$$

Note that for $i < l$,

$$\begin{aligned} \Delta p(b_i) &= \left(\frac{1}{2}\right)^{n-1} \left[\frac{(n-1)!}{(i-1)!(n-i)!} - 8\frac{(n-4)!}{(i-1)!(n-i-3)!}\right] \\ &= \left(\frac{1}{2}\right)^{n-1} \frac{(n-4)!}{(i-1)!(n-i)!} \underbrace{[(n-1)(n-2)(n-3) - 8(n-i)(n-i-1)(n-i-2)]}_{f(i)} \end{aligned}$$

Since $f(i)$ is monotonically increasing in i, then by observing that $f\left(\frac{n}{2} - 1\right) < 0$ and $f\left(\frac{n}{2}\right) > 0$, we have

$$\Delta p(b_i) \text{ is } \begin{cases} < 0 \text{ if } i < \min(l, n/2) \\ > 0 \text{ if } n/2 \le i < l \end{cases}$$

Similarly, for $i > l+2$, we can use an analogous argument (or apply symmetry) to claim that,

$$\Delta p(b_i) \text{ is } \begin{cases} < 0 \text{ if } i > \max(l+2, n/2) \\ > 0 \text{ if } l+2 < i \leq n/2 \end{cases}$$

Finally, it is clear that $\Delta p(b_i) > 0$ for $i = l, l+1, l+2$, so we are done. □

LEMMA B.2. *Let R and S be urns as defined in Section 5.1.1. From R, create a new urn R' by moving all balls left of x to x and all balls right of z to z. Similarly, from S, create a new urn S' by moving all balls left of x to x and all balls right of z to z. Let $\Delta U'_x = U_x^{S'} - U_x^{R'}$ and $\Delta U'_z = U_z^{S'} - U_z^{R'}$. Then,*

$$\Delta U_x \geq \Delta U'_x \quad \textit{and} \quad \Delta U_z \geq \Delta U'_z$$

PROOF. Since the relative positions of the balls have not changed in R' and S', the change in winning probabilities from R' to S' are the same, i.e. $\Delta p(b_i) = \Delta p'(b_i)$. Then for $\Delta p(b_i) < 0$,

$$\begin{aligned} \Delta p(b_i) U_{x(z)}(b_i) \geq \Delta p'(b_i) U'_{x(z)}(b_i) &\iff U_{x(z)}(b_i) \leq U'_{x(z)}(b_i) \\ &\iff b_i \text{ is closer to } x(z) \end{aligned}$$

and for $\Delta p(b_i) > 0$,

$$\begin{aligned} \Delta p(b_i) U_{x(z)}(b_i) \geq \Delta p'(b_i) U'_{x(z)}(b_i) &\iff U_{x(z)}(b_i) \geq U'_{x(z)}(b_i) \\ &\iff b_i \text{ is farther from } x(z) \end{aligned}$$

Then we can use this to prove our Lemma for each of three cases:

Case 1: $b_{n/2} \in [x, z]$ for urn S. When the median ball is located in $[x, z]$, all the balls left of x and right of z satisfy $\Delta p(b_i) < 0$. Then moving the balls left of x to x brings them closer to both x and z. Similarly, moving the balls right of z to z brings them closer to both x and z. Therefore, we must have

$$\Delta U_{x(z)} = \sum_{i=1}^{n} \Delta p(b_i) U_{x(z)}(b_i) \geq \sum_{i=1}^{n} \Delta p'(b_i) U'_{x(z)} = \Delta U'_{x(z)}$$

Case 2: $b_{n/2} < x$ for urn S. When the median ball is located to the left of x, we will need to first make an intermediate pair of urns. Note that the balls left of $b_{n/2}$ have $\Delta p(b_i) < 0$ so we can move them rightwards to bring them closer to both x and z. However, since the balls between $b_{n/2}$ and x have $\Delta p(b_i) > 0$, we need to move them leftwards to bring them farther away from both x and z. For the balls right of z, $\Delta p(b_i) < 0$, so we can again move them to z which brings them closer to x and z. Then let urns R'', S'' be another pair of urns where balls left of x are moved to the position of ball $b_{n/2}$ and balls right of z are moved to z. By an argument similar to that of Case 1, $\Delta U_{x(z)} \geq \Delta U''_{x(z)}$. Now, note that

$$\Delta U''_{x(z)} = \left[\sum_{i=1}^{k-1} \Delta p''(b_i) \right] U''_{x(z)}(b_{n/2}) + \sum_{i=k}^{n} \Delta p''(b_i) U''_{x(z)}(b_i)$$

where k is the index of the leftmost ball that is right of x and U'', p'' are the analogous expressions for utility and winning probability in urns R'', S''. But we also know that for $k < l$,

$$\sum_{i=1}^{k-1} \Delta p''(b_i) = \sum_{i=1}^{k-1} \left[\left(\frac{1}{2}\right)^{n-1} \binom{n-1}{i-1} - \left(\frac{1}{2}\right)^{n-4} \binom{n-4}{i-1} \right]$$
$$= \mathbf{Pr}[\leq k-2 \text{ heads in } n-1 \text{ coin flips}] - \mathbf{Pr}[\leq k-2 \text{ heads in } n-4 \text{ coin flips}] \leq 0$$

Then this means that

$$\left[\sum_{i=1}^{k-1} \Delta p''(b_i)\right] U''_{x(z)}(b_{n/2}) \geq \left[\sum_{i=1}^{k-1} \Delta p'(b_i)\right] U'_{x(z)}(b_{n/2}) \iff U''_{x(z)}(b_{n/2} \leq U'_{x(z)}(b_{n/2})$$
$$\iff b_{n/2} \text{ is closer to } x(z)$$

We know that by moving the balls at the position of ball $b_{n/2}$ rightwards to x, we are bringing it closer to both x and z to get R', S'. But we just showed that this decreases $\Delta U''_{x(z)}$, so $\Delta U''_{x(z)} \geq U'_{x(z)}$. This finishes the proof for this case.

Case 3: $b_{n/2} > z$ for urn S. The proof is symmetric to that of Case 2. □

THEOREM B.3. *Let R and S be urns as defined in Section 5.1.1 and let all their balls lie within the interval $[x, z]$. Then if participants x and z have concave utility functions, at least one of $\Delta U_x \geq 0$ and $\Delta U_z \geq 0$ is true.*

PROOF. Let $A = \min(l, n/2)$ and $Z = \max(l+2, n/2)$ be the leftmost and rightmost balls for which $\Delta p(b_i) > 0$. Recall that by Lemma 5.1, all balls left of A and right of Z must have $\Delta p(b_i) < 0$. Then we can separate the expression for ΔU_x into three summations,

$$\Delta U_x = \sum_{i=1}^{A-1} \Delta p(b_i) U_x(b_i) + \sum_{i=A}^{Z} \Delta p(b_i) U_x(b_i) + \sum_{i=Z+1}^{n} \Delta p(b_i) U_x(b_i)$$

Note that $\sum_{i=1}^{A-1} \Delta p(b_i) + \sum_{i=Z+1}^{n} \Delta p(b_i) = -\sum_{i=A}^{Z} \Delta p(b_i)$. Then we can partition up the $\Delta p(b_j)$ mass for $j \in [A, Z]$ among the balls $b_i, i \notin [A, Z]$ such that the mass assigned to a given ball exactly matches the magnitude of $\Delta p(b_i)$. Since the $\Delta p(b_j)$ values for $j \in [A, Z]$ may not exactly match the values for the balls $i \notin [A, Z]$, we may need to use several of the $\Delta p(b_j)$ terms from $j \in [A, Z]$ for any given index. Represent such a mapping with $U_x^*(i)$. Then we can factor this to get

$$\Delta U_x = \sum_{i=1}^{A-1} \Delta p(b_i)[U_x(b_i) - U_x^*(i)] + \sum_{i=Z+1}^{n} \Delta p(b_i)[U_x(b_i) - U_x^*(i)]$$

where $U_x^*(i)$ is some convex combination of utilities $U_x(b_j), j \in [A, Z]$. Note that since balls are indexed left to right and all lie within $[x, z]$, then balls indexed $i \in [1, A)$ are closer to x than those indexed $i \in [A, Z]$, which are closer than those indexed $i \in (R, n]$. Therefore, we have

$$U_x(b_i) - U_x^*(i) \text{ is } \begin{cases} \geq 0 & \text{if } i = 1, 2, \ldots A-1 \\ \leq 0 & \text{if } i = Z+1, Z+2, \ldots, n \end{cases}$$

Combining these, we get that

$$U_x \geq 0 \iff \frac{\sum_{i=Z+1}^{n} \Delta p(b_i)[U_x(b_i) - U_x^*(i)]}{\sum_{i=1}^{A-1} \Delta p(b_i)[U_x^*(i) - U_x(b_i)]} \geq 1 \tag{1}$$

since $\Delta p(b_i)[U_x^*(i) - U_x(b_i)] \geq 0$ for $i = 1, 2, \ldots A-1$. Similarly, we have

$$\Delta U_z = \sum_{i=1}^{A-1} \Delta p(b_i)[U_z(b_i) - U_z^*(i)] + \sum_{i=Z+1}^{n} \Delta p(b_i)[U_z(b_i) - U_z^*(i)]$$

where $U_z^*(i)$ is the convex combination of utilities $U_z(b_j), j \in [A, Z]$ that uses the same indices j as those used in $U_x^*(i)$. Now, for z, balls indexed $i \in [1, A)$ are farther from z than those indexed $i \in [A, Z]$, which are farther from those indexed $i \in (Z, n]$. Therefore

$$U_z(b_i) - U_z^*(i) \text{ is } \begin{cases} \leq 0 & \text{if } i = 1, 2, \ldots A-1 \\ \geq 0 & \text{if } i = Z+1, Z+2, \ldots, n \end{cases}$$

Combining these, we get that

$$U_z \geq 0 \iff \frac{\sum_{i=Z+1}^{n} \Delta p(b_i)[U_z(b_i) - U_z^*(i)]}{\sum_{i=1}^{A-1} \Delta p(b_i)[U_z^*(i) - U_z(b_i)]} \leq 1 \tag{2}$$

since $\Delta p(b_i)[U_z^*(i) - U_z(b_i)] \leq 0$ for $i = 1, 2, \ldots A-1$. We now have one last step. For any $f(x)$ which is concave and monotonically decreasing, we have that

$$\frac{\sum_{i=1}^{m} c_i[f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^{n} d_j[f(s_j^2) - f(s_j^1)]} \geq \frac{\sum_{i=1}^{m} c_i[t_i^2 - t_i^1]}{\sum_{j=1}^{n} d_j[s_j^2 - s_j^1]} \quad \text{and} \quad \frac{\sum_{i=1}^{m} c_i[f(s_i^1) - f(s_i^2)]}{\sum_{j=1}^{n} d_j[f(t_j^1) - f(t_j^2)]} \leq \frac{\sum_{i=1}^{m} c_i[s_i^1 - s_i^2]}{\sum_{j=1}^{n} d_j[t_j^1 - t_j^2]}$$

for $s_j^1 \leq t_i^1$, $s_j^2 \leq t_i^2$, $s_i^1 \leq s_j^2$, $t_i^1 \leq t_i^2$, and $\text{sign}(c_i) = \text{sign}(d_j)$ (as detailed in Appendix C). Applying this to (1) and (2), we get

$$\frac{\sum_{i=Z+1}^{n} \Delta p(b_i)[U_x(b_i) - U_x^*(i)]}{\sum_{i=1}^{A-1} \Delta p(b_i)[U_x^*(i) - U_x(b_i)]} \geq \frac{\sum_{i=Z+1}^{n} \Delta p(b_i)[d(b_i, b^*(i))]}{\sum_{i=1}^{L-1} \Delta p(b_i)[d(b^*(i), b_i)]}$$

and

$$\frac{\sum_{i=Z+1}^{n} \Delta p(b_i)[U_z(b_i) - U_z^*(i)]}{\sum_{i=1}^{A-1} \Delta p(b_i)[U_z^*(i) - U_z(b_i)]} \leq \frac{\sum_{i=Z+1}^{n} \Delta p(b_i)[d(b_i, b^*(i)]}{\sum_{i=1}^{A-1} \Delta p(b_i)[d(b^*(i), b_i)]}$$

where $d(b_i, b^*(i))$ is a convex combination of $d(b_i, b_j)$ for $j \in [A, Z]$, where the weights are the same as those of $U_{x(z)}^*$. Then either

$$\frac{\sum_{i=Z+1}^{n} \Delta p(b_i)[d(b_i, b^*(i)]}{\sum_{i=1}^{A-1} \Delta p(b_i)[d(b^*(i), b_i)]} \geq 1 \quad \text{or} \quad \frac{\sum_{i=Z+1}^{n} \Delta p(b_i)[d(b_i, b^*(i)]}{\sum_{i=1}^{A-1} \Delta p(b_i)[d(b^*(i), b_i)]} \leq 1$$

If the first is true, we can apply (1) to claim that $U_x \geq 0$. If the second is true, we can apply (2) to claim that $U_z \geq 0$. Since one of these must be true, we are done. □

C. SUPPORTING PROOFS

THEOREM C.1. *[Lee and Bruck 2012] Let a fixed size urn with R_0 red balls out of n total balls have an urn function $f(p)$ for which $\frac{f(p)}{f(1-p)}$ is monotonically decreasing and let T denote the first time when either $R_T = n$ or $R_T = 0$. Then,*

$$\mathbb{E}[T] \leq \frac{1}{q_1} + \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor - 1} \frac{q_k}{q_{k+1}(q_k - p_k)}$$

where $p_k = f\left(\frac{n-k}{n}\right)$ and $q_k = f\left(\frac{k}{n}\right)$.

COROLLARY C.2. *Let a fixed size urn with R_0 red balls out of n total balls have an urn function $f(p) = 3p(1-p)^2$ and let T denote the first time when either $R_T = n$ or $R_T = 0$. Then,*

$$\mathbb{E}[\tau] \leq n \ln n + O(n)$$

PROOF. From Theorem C.1, we know that

$$\mathbb{E}[T] \leq \frac{1}{q_1} + \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor - 1} \frac{q_k}{q_{k+1}(q_k - p_k)} = \frac{n^3}{3(n-1)^2} + \frac{n^3}{3} \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor - 1} \underbrace{\frac{n-k}{(k+1)(n-k-1)^2(n-2k)}}_{(*)}$$

where

$$\begin{aligned}
(*) &= -\frac{1}{n(n-2)(n-k-1)^2} + \frac{4n}{(n+2)(n-2)^2(n-2k)} \\
&\quad + \frac{n+1}{n^2(n+2)(k+1)} - \frac{n^2+n-2}{n^2(n-2)^2(n-k-1)} \\
&\leq \frac{4n}{(n+2)(n-2)^2(n-2k)} + \frac{n+1}{n^2(n+2)(k+1)} \\
&= \left(\frac{4}{n^2(n-2k)} + \frac{1}{n^2(k+1)} \right) \left(1 + O\left(\frac{1}{n}\right) \right)
\end{aligned}$$

Putting it together, we have

$$\mathbb{E}[T] \leq \frac{n}{3} + o(1) + \frac{n}{3} \left(1 + O\left(\frac{1}{n}\right) \right) \sum_{k=1}^{\lfloor \frac{n}{2} \rfloor - 1} \left[\frac{4}{n-2k} + \frac{1}{k+1} \right]$$

Noting that $\sum_{i=1}^{k} \frac{1}{k} = H_k = \ln n + O(1)$, where H_k is the k-th Harmonic number, we have

$$\begin{aligned}
\mathbb{E}[T] &\leq O(n) + \frac{n}{3} \left(2 \ln \frac{n}{2} + \ln \frac{n}{2} \right) \\
&= n \ln n + O(n)
\end{aligned}$$

□

LEMMA C.3. *Let $f(x)$ be a concave monotonically decreasing function. Then for $s_j^1, s_j^2, t_i^1, t_i^2 \in \mathbb{R}$ satisfying $s_j^1 \leq t_i^1$, $s_j^2 \leq t_i^2$, $s_i^1 \leq s_j^2$, and $t_i^1 \leq t_i^2$ and sign(c_i) = sign(d_j), we have*

$$\frac{\sum_{i=1}^{m} c_i [f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^{n} d_j [f(s_j^2) - f(s_j^1)]} \geq \frac{\sum_{i=1}^{m} c_i [t_i^2 - t_i^1]}{\sum_{j=1}^{n} d_j [s_j^2 - s_j^1]} \quad \textit{and} \quad \frac{\sum_{i=1}^{m} c_i [f(s_i^1) - f(s_i^2)]}{\sum_{j=1}^{n} d_j [f(t_j^1) - f(t_j^2)]} \leq \frac{\sum_{i=1}^{m} c_i [s_i^1 - s_i^2]}{\sum_{j=1}^{n} d_j [t_j^1 - t_j^2]}$$

PROOF. Since f is concave, $s_j^1 \leq t_i^1$, and $s_j^2 \leq t_i^2$,

$$\frac{f(t_i^2) - f(t_i^1)}{t_i^2 - t_i^1} \leq \frac{f(t_i^2) - f(s_j^1)}{t_i^2 - s_j^1} \leq \frac{f(s_j^2) - f(s_j^1)}{s_j^2 - s_j^1}$$

Then by noting that f is monotonically decreasing, $s_j^1 \leq s_j^2$, $t_i^1 \leq t_i^2$, and $\text{sign}(c_i) = \text{sign}(d_j)$, we achieve the statement for a single term on the top and bottom

$$\frac{c_i[f(t_i^2) - f(t_i^1)]}{d_j[f(s_j^2) - f(s_j^1)]} \geq \frac{c_i[t_i^2 - t_i^1]}{d_j[s_j^2 - s_j^1]}$$

If we flip this inequality, we get

$$\frac{d_j[f(s_j^2) - f(s_j^1)]}{c_i[f(t_i^2) - f(t_i^1)]} \leq \frac{d_j[s_j^2 - s_j^1]}{c_i[t_i^2 - t_i^1]}$$

Using this, we can derive,

$$\frac{\sum_{j=1}^{n} d_j[f(s_j^2) - f(s_j^1)]}{c_i[f(t_i^2) - f(t_i^1)]} = \sum_{j=1}^{n} \frac{d_j[f(s_j^2) - f(s_j^1)]}{c_i[f(t_i^2) - f(t_i^1)]} \leq \sum_{j=1}^{n} \frac{d_j[s_j^2 - s_j^1]}{c_i[t_i^2 - t_i^1]} = \frac{\sum_{j=1}^{n} d_j[s_j^2 - s_j^1]}{c_i[t_i^2 - t_i^1]}$$

Then by flipping this inequality, we get

$$\frac{c_i[f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^{n} d_j[f(s_j^2) - f(s_j^1)]} \geq \frac{c_i[t_i^2 - t_i^1]}{\sum_{j=1}^{n} d_j[s_j^2 - s_j^1]}$$

Finally, we can use this to derive the first part of our final result

$$\frac{\sum_{i=1}^{m} c_i[f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^{n} d_j[f(s_j^2) - f(s_j^1)]} = \sum_{i=1}^{m} \frac{c_i[f(t_i^2) - f(t_i^1)]}{\sum_{j=1}^{n} d_j[f(s_j^2) - f(s_j^1)]} \geq \sum_{i=1}^{m} \frac{c_i[t_i^2 - t_i^1]}{\sum_{j=1}^{n} d_j[s_j^2 - s_j^1]} = \frac{\sum_{i=1}^{m} c_i[t_i^2 - t_i^1]}{\sum_{j=1}^{n} d_j[s_j^2 - s_j^1]}$$

We can get the second part by simply inverting and multiplying the top and bottom of both sides by -1

$$\frac{\sum_{j=1}^{n} d_j[f(s_j^2) - f(s_j^1)]}{\sum_{i=1}^{m} c_i[f(t_i^2) - f(t_i^1)]} \leq \frac{\sum_{j=1}^{n} d_j[s_j^2 - s_j^1]}{\sum_{i=1}^{m} c_i[t_i^2 - t_i^1]}$$

□