Thwarting Vote Buying Through Decoy Ballots
Extended Version

David C. Parkes[1], Paul Tylkin[1(✉)], and Lirong Xia[2]

[1] Harvard University, Cambridge, MA, USA
ptylkin@g.harvard.edu
[2] Rensselaer Polytechnic Institute, Troy, NY, USA

Abstract. There is increasing interest in promoting participatory democracy, in particular by allowing voting by mail or internet and through random-sample elections. A pernicious concern, though, is that of *vote buying*, which occurs when a bad actor seeks to buy ballots, paying someone to vote against their own intent. This becomes possible whenever a voter is able to sell evidence of which way she voted. We show how to thwart vote buying through *decoy ballots*, which are not counted but are indistinguishable from real ballots to a buyer. We show that an Election Authority can significantly reduce the power of vote buying through a small number of optimally distributed decoys, and model societal processes by which decoys could be distributed. We also introduce a generalization of our model to non-binary election outcomes.

1 Introduction

The goal of participatory democracy [9,11] is to engage citizens more frequently and with more granularity in the decision-making processes of government bodies. Technologies that can help with this transition are those that support voting from the home by mail or over the internet, and that make use of *random sample elections*, in which a representative subsample of the population is tasked with voting on a particular issue, allowing participatory democracy to function without everyone needing to be concerned with every issue.

A pernicious concern, though, is that of *vote buying*, where a bad actor attempts to gain improper influence in an election by purchasing ballots from voters and paying them to vote against their intent. The practical implications of this are manifold, since the social construct of elections relies on the perception of reliability and fairness. Vote buying has been an everlasting threat to democracy; for example, a survey shows that in the 1996 Thai general elections

This is an extended version of *Thwarting Vote Buying through Decoy Ballots*, which was presented at the EXPLORE 2017 workshop at the 16th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2017) and the Proc. 26th Int. Joint Conf. on Artificial Intelligence (IJCAI 2017).

G. Sukthankar and J. A. Rodriguez-Aguilar (Eds.): AAMAS 2017 Visionary Papers, LNAI 10643, pp. 45–66, 2017.
https://doi.org/10.1007/978-3-319-71679-4_4

"*one third of households were offered money to buy votes at the last general election*" [13]. Schaffer [14] mentions that "*[Vote buying]... is making an impressive comeback...it seems, a blossoming market for votes has emerged as an epiphenomenon of democratization*". New technologies can make the situation worse. For example, web platforms can serve as middlemen, digital currency supports anonymous payments, and abundant data coupled with machine learning can help buyers discover entrapment schemes as well as identify voters to target with offers.

In this paper, we show that vote buying can be thwarted by distributing *decoy ballots*, which are not counted, in addition to real ballots. A vote buyer will not know whether a ballot is real or decoy, and thus, decoys (if sold) may deplete a buyer's budget. Voters who know that they have a decoy ballot are motivated to sell their ballots to a buyer, both for reasons of profit and out of civic duty, wanting to maintain the election's integrity. David Chaum earlier introduced the notion of random sample voting, and proposed decoy ballots in order to address the potential problem of vote buying in remote elections generally and for random sample voting in particular [4]. He has also introduced the key notion of *proof of decoy* (see Sect. 2). We study how to distribute decoy ballots, and analyze the power of this approach.

We assume that real ballots impose a high cost on society, for the reason that it takes effort for citizens to become informed about an issue and vote appropriately, thus representing their considered opinion on an issue.[1] Without the willingness to invest this effort, methods of participatory democracy may ultimately fail. For example, a simple calculation for the US shows that if we assume that 200 M people will participate, and there are about 12,000 issues to decide per year,[2] then assuming that voters are willing to engage three times a year, we have a maximum of 50,000 voters per issue. At this scale, vote buying, especially on contentious issues, may pose a severe problem.

Turning to decoy ballots, we model these as costly but not so costly that the number of decoys to distribute cannot be considered as a design decision of the Election Authority. The cost of decoys comes about because, to be effective, voters need to be willing to go to the effort to sell the ballot (and thus, cast the ballot and prove which way it was cast) if approached by a buyer. But because any decoy ballots are not counted, we assume it is less cognitively expensive for a voter to form an opinion.

Although we situate our discussion in a societal context, similar themes can be imagined for economies of AIs [12], where it is desired to elicit and fairly aggregate multiple opinions, but would not be scalable to request input from every agent all the time.

Our Contributions. Focusing mostly on the binary outcome case, we provide a formal model of vote buying, including a characterization of the vote buyer's

[1] In some approaches to random-sample voting this cost comes also about as a result of needing to physically mail ballots.

[2] This represents the approximate voter population and the number of issues before Congress per year, assuming 2 issues per bill.

behavior and an optimal policy for distributing decoy ballots by the Election Authority (EA). In addition, we model two societal processes by which decoys could be distributed—these approaches freeing the EA of any concern that it could be seen to be biasing the outcome of an election when distributing decoys in any way other than reflecting a random sample of the population. In simulation, we show that the EA can make effective use of decoy ballots to maintain election integrity (e.g., reducing the probability that the buyer changes the outcome to less than 1%). For the optimal defense, we are able to achieve this by adding a small number of decoys that are proportional in quantity to the number of ballots the buyer can afford to buy. Interestingly, a "civic duty defense" that allocates decoys to a random subset of those who request one is almost as effective as the optimal defense in which the EA optimizes the distribution of voter types that receive decoys. We also provide a generalization of our model to the three-outcome case, prove that a *buy the expected winner* strategy is optimal for elections with simple voter types, and provide numerical results illustrating the strategy of both the buyer and the EA in equilibrium.

Related Work. There are numerous studies on vote buying, for example [8, 15, 16, 19]. These include game-theoretic models of vote buying, but none that consider the role of decoy ballots. In the work by Dekel *et al.* [6], the game is played by the candidates themselves buying votes, Groseclose and Snyder [10] study vote buying in legislative bodies and analyze the optimal coalition size. Vicente [18] studies the incumbency advantage in a vote buying game. Within AI, the problem studied here related to studies of control (manipulation of the election structure, including changing the candidate slate) and bribery (voters are paid by an interested party to vote a certain way) as studied in computational social choice [2, 7]. In particular, the *lobbying* problem considers an election with a binary outcome on a number of issues, and the vote buyer has a total budget that can be expended across all issues [1, 3, 5]. Ours is a special case with a single issue, but whereas previous research has focused on using computational complexity as a barrier against bribery and control, we adopt a game-theoretic model and study the power of decoy ballots. There is also a conceptual connection with work on *security games* [17], where the approach is to use game theory to design optimal strategies to prevent losses from terrorist attacks.

2 The Model

We assume that there is a large population of possible voters, and, for now, assume that this is a binary choice election with possible votes YES and NO. For expositional simplicity, we assume that all voters who receive a real ballot will place a vote. Similarly, we assume that every voter for whom it is profitable to sell a ballot (decoy or otherwise) will try to sell the ballot.[3]

[3] It is simple to generalize the model so that the voters who cast ballots are sampled uniformly from those who receive ballots, and similarly for those who try to sell ballots.

The voters. Each voter i has an immutable, publicly-observable *voter type*, θ_i, which indicates the probability that a random voter with this type will vote YES. We can think about θ_i as the prior that a voter will vote YES before she has carefully considered the merits of an issue. Voter types are drawn independently from a *voter type distribution* with probability density f, assumed to have full support on $[0, 1]$. We assume without loss of generality that $E_f[\theta] < 1/2$, i.e., that the outcome of the election without any interference by a buyer and with enough real ballots is NO.

The buyer. We model a single, budget-limited buyer. Given our assumption that $E_f[\theta] < 1/2$, we consider the interesting case of a *YES-buyer*, meaning that the buyer wants the election outcome to be YES. To keep things simple, we assume the buyer can find the voters with ballots, and will offer the same price $p > 0$ to each voter in some subset of these voters. The buyer has a budget B, representing the number of ballots that he can afford to purchase at price p, and has no utility for unspent budget. The buyer selects a random subset of voters if more respond to the offer than he can afford.

Conditioned on whether a voter's intent is to vote NO or YES, and whether they have a real or decoy ballot, all voters have the same utility function in regard to whether or not to sell. In particular, simple analysis yields that this ordering of the minimum price that a voter will require in order to agree to sell a ballot is real-NO > real-YES > decoy-YES > decoy-NO. For example, any price that is acceptable to a "real-YES" voter (real ballot, intent to vote YES) is also acceptable to "decoy-YES" and "decoy-NO" voters. Ballots from decoy-NO voters are the cheapest to buy.[4]

Based on this, the real-NO votes—and the only ones the buyer is interested in—are the most expensive ballots to buy. Because of this, we assume the buyer will set price p high enough for a real-NO voter to agree to sell if approached. This could be set based on market research, for example.

The game form. The voters who receive a real ballot are a random subset of the population, and thus with types that follow f. The choice of how to distribute decoy ballots is, in general, a design decision. Let ψ denote the density function for this decoy ballot distribution. Modeled as a sequential-move game, the election proceeds in three stages:

(1) The EA distributes some number of real and decoy ballots, with the number and type distribution of real ballots assumed fixed, but the number of decoy ballots, and perhaps type distribution ψ a design decision.

[4] To understand this ordering, suppose that a voter with a real ballot has a cost for selling, representing the possibility of being caught. In addition, voters that intend to vote NO prefer not to vote YES. Thus, real-NO ballots are the most expensive votes to buy. Amongst decoys, decoy-YES ballots are more expensive to buy than decoy-NO ballots because a voter who would vote NO (if she had a real ballot) has a value for depleting the budget of a YES-buyer. This is not the case for a voter who would vote YES.

(2) The buyer learns who has received a ballot (possibly a decoy) and chooses to offer price p to each voter in some subset of voters who have (real or decoy) ballots. The voters who receive an offer decide whether or not to sell. The buyer breaks ties at random if multiple voters agree to sell.
(3) Both real and decoy ballots are cast, and the real ballots are tallied to determine the outcome. The buyer makes payments to voters who agreed to sell and provide a proof that they vote YES.

Both distribution f and the type of each voter is common knowledge. Our analysis will focus on the subgame perfect equilibrium of this game. Throughout, the voters have a simple equilibrium behavior—agree to sell if offered a price p (which will, in equilibrium, be high enough to be acceptable.)

Proof of decoy. We assume the existence of a *proof-of-decoy*, which lets a voter with a decoy choose to prove that she has a decoy. This is required to mitigate the "fear of being caught selling a ballot"— that way, a voter with a decoy can prove to a vigilante that she is not selling a real ballot. On the other hand, there is no way to prove the authenticity of a real ballot. This property is easy to support through standard cryptographic primitives; see, for example, Chaum [4].[5]

EA and Buyer objectives. We take as the objective of the EA that of maintaining election integrity, and thus *minimizing the probability that the buyer changes the election outcome*. In contrast, the interests of the buyer are diametrically opposed, and he wants to *maximize the probability that the outcome of the election is changed*.

3 Buyer Analysis

Given the buyer's objective, the best response of the buyer to the EA is to maximize the expected number of real-NO ballots that he buys, given his budget B and knowledge about voters' types (probability of voting YES). Let $\mathcal{I} \subseteq [0, 1]$ denote the subset of voter types from which the buyer buys; in particular, the buyer will buy every ballot held (real or decoy) by voters of these types. Let n_r denote the number of real ballots and n_d the number of decoy ballots. The buyer wants to select the subset $\mathcal{I}$ to solve:

$$\max_{\mathcal{I}} \int_{\mathcal{I}} \frac{n_r}{n_r + n_d}(1-\theta) f(\theta) d\theta \text{ s.t. } \int_{\mathcal{I}} n_r f(\theta) + n_d \psi(\theta) d\theta \le B. \tag{1}$$

In this way, the buyer maximizes a quantity that is proportional to the expected number of real-NO ballots purchased, subject to the total budget.

[5] The asymmetry in having proof-of-decoy without proof-of-authenticity is important to prevent a buyer from using coercion to buy only real ballots, while at the same time allowing a voter with a decoy ballot to sell without fear of being accused of acting against the social good. A voter will never choose to reveal that she holds a decoy to a buyer, since doing so would remove the chance of a sale.

Let $h(\theta)$ denote the *probability that a ballot is real-NO given type θ*. By Bayes' rule, and recalling that the buyer has knowledge of f and ψ, this is

$$h(\theta) \stackrel{\text{def}}{=} P(\text{real} \wedge \text{NO}|\theta) = \frac{n_r(1-\theta)f(\theta)}{n_r f(\theta) + n_d \psi(\theta)}. \tag{2}$$

Given a set $I \subseteq [0,1]$, let $h(I)$ denote the set $\{h(\theta)\}$ for $\theta \in I$. Let $h(I_1) < h(I_2)$ mean that every value in I_1 is strictly less than every value in I_2.

Lemma 1 (Buyer Optimality). *The optimal buyer strategy in the subgame perfect equilibrium is to buy in order of decreasing $h(\theta)$ until the budget is exhausted.*

Proof. Suppose not, i.e., suppose that there is a set $J \subset \mathcal{I}$ and a set $J' \not\subset \mathcal{I}$ such that $h(J') > h(J)$. Then, the buyer could strictly increase his objective by buying J' instead of J.

We assume w.l.o.g. that if a YES-buyer has to choose between buying two subsets of $[0,1]$ for which $h(\theta)$ is equal, he will buy the subset with lower θ. Let $\mathfrak{M} \stackrel{\text{def}}{=} \int_{\mathcal{I}} f(\theta)d\theta$ denote the fraction of real ballots that the buyer buys. By 'election bought', we refer to the event that the buyer buys enough real ballots to change the outcome (with n_r real ballots); by 'correct outcome is NO', we refer to the event that the election outcome is NO (with $n_r + n_d$ real ballots).

Lemma 2. *The probability that the buyer changes the outcome in the subgame perfect equilibrium is given by*

$$\begin{aligned} & P(\textit{buyer changes outcome}) \\ = \; & P([\textit{election bought}] \wedge [\textit{correct outcome is NO}]) \\ \approx & P\left(\frac{n_r(1-2\mathfrak{M}-2(1-\mathfrak{M})\mu_Y)}{2\sqrt{n_r(1-\mathfrak{M})\mu_Y(1-\mu_Y)}} < Z < \frac{(1-2\mu)\sqrt{n_r+n_d}}{2\sqrt{\mu(1-\mu)}}\right), \end{aligned} \tag{3}$$

where $Z \sim \mathcal{N}(0,1)$, $\mu \stackrel{\text{def}}{=} E_f[\theta]$, and $\mu_Y \stackrel{\text{def}}{=} \frac{1}{1-\mathfrak{M}} \int_{[0,1]\setminus\mathcal{I}} \theta f(\theta) d\theta$.

Proof. Let the type distribution of the unbought types be given by

$$f_Y(\theta) \stackrel{\text{def}}{=} \begin{cases} \frac{f(\theta)}{1-\mathfrak{M}} & \text{for } \theta \in [0,1] - \mathcal{I} \\ 0 & \text{for } \theta \in \mathcal{I} \end{cases}. \tag{4}$$

To model votes, we introduce the shorthand notation $X_i \multimap f(\theta)$ to denote the hierarchical model $\theta_i \sim f(\theta); X_i \sim \text{Bern}(\theta_i)$. The probability that the buyer changes the outcome is given by

$$\begin{aligned} & P(\textit{buyer changes outcome}) \\ = \; & P([\textit{election bought}] \wedge [\textit{correct outcome is NO}]) \\ = \; & P\left(\left[\frac{\sum_{i=1}^{(1-\mathfrak{M})n_r} V_i}{n_r} + \mathfrak{M} > \frac{1}{2}\right] \wedge \left[\frac{\sum_{j=1}^{n_r+n_d} W_j}{n_r+n_d} < \frac{1}{2}\right]\right), \end{aligned} \tag{5}$$

where $V_i \multimap f_Y(\theta)$ and $W_j \multimap f(\theta)$. We can use the Normal approximation to the Binomial to obtain

$$P(\textit{buyer changes outcome})$$
$$\approx P\left(\frac{n_r(1-2\mathfrak{M}-2(1-\mathfrak{M})\mu_Y)}{2\sqrt{n_r(1-\mathfrak{M})\mu_Y(1-\mu_Y)}} < Z < \frac{(1-2\mu)\sqrt{n_r+n_d}}{2\sqrt{\mu(1-\mu)}}\right). \tag{6}$$

This allows us to compute the probability the buyer changes the election outcome, which is determined by the fraction of real ballots that he is able to buy given a defense.

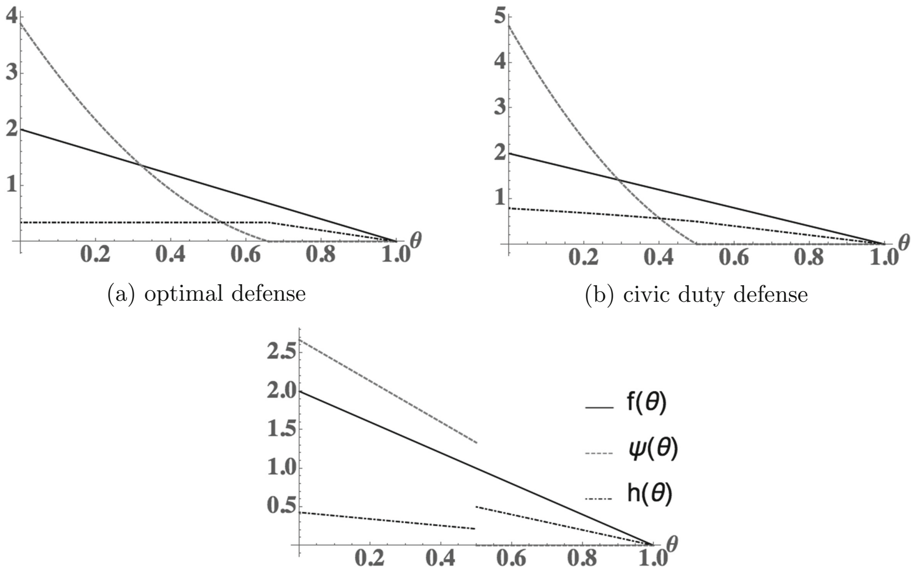

(a) optimal defense

(b) civic duty defense

(c) auction-based defense

Fig. 1. Examples of type distribution $f(\theta)$, decoy distribution $\psi(\theta)$, and desirability to buyer $h(\theta)$ for (a) an optimal defense, (b) a civic duty defense with max type requesting a decoy $x_C = 0.5$ and 10% decoy ballots, (c) an auction-based defense with max type assigned a decoy $x_A = 0.5$ and 50% decoy ballots. Here $f = \text{Beta}(1, 2)$.

4 Optimal Decoy Distribution

In this section, we assume that the EA can design defense distribution ψ, and study the equilibrium of the vote-buying game where the EA chooses an optimal defense given that the buyer will best respond.

Definition 1 (Canonical Defense). *Defense ψ is* canonical *if there is some x, $0 \leq x \leq 1$, s.t. $h(\theta) = \min(1-x, 1-\theta)$.*

See Fig. 1(a) for an illustration of a canonical defense. Let supp(g) denote the support of distribution g. Define the following two properties for ψ:

(P1) $h(\theta)$ has the same value for all $\theta \in \text{supp}(\psi)$.
(P2) $\min_{\theta \in \text{supp}(\psi)} h(\theta) \geq \max_{\theta \notin \text{supp}(\psi)} h(\theta)$

Lemma 3. *Any defense ψ satisfying both P1 and P2 is canonical.*

Proof. We assume that ψ satisfies P1 and P2, and show that $\text{supp}(\psi) = [0, \text{xx}_\text{o}]$ for some $x_\text{o} \in [0, 1]$, i.e., we must have $0 \in \text{supp}(\psi)$, supp(ψ) must be contiguous, and the left endpoint of $\text{supp}(\psi) \stackrel{\text{def}}{=} [x_0, x_\text{o}]$ is 0 (i.e., x_0 must be 0). Assume that ψ is a defense that satisfies both P1 and P2.

Since $\forall \theta \notin \text{supp}(\psi), h(\theta) = 1 - \theta$ and $\forall \theta \in \text{supp}(\psi), h(\theta) \leq 1 - \theta$, this tells us that P2 requires $0 \in \text{supp}(\psi)$. Otherwise, $h(0) > \max_{\theta \in \text{supp}(\psi)} h(\theta) \geq \min_{\theta \in \text{supp}(\psi)} h(\theta)$, which contradicts P2.

Next, assume for contradiction that supp(ψ) is not contiguous. Then, consider the first two intervals $J_1 \stackrel{\text{def}}{=} [x_1, x_2]$ and $J_2 \stackrel{\text{def}}{=} [x_3, x_4]$, with $J_1, J_2 \subseteq \text{supp}(\psi)$. By P1, $h(\theta)$ has the same value $\forall \theta \in \text{supp}(\psi)$. Call this value y. First, we examine the special case of $x_1 = 0$. Then, we have

$$y \leq (1 - x4) < (1 - x3) \leq \max_{\theta \notin \text{supp}(\psi)} h(\theta), \tag{7}$$

i.e., $y < \max_{\theta \notin \text{supp}(\psi)} h(\theta)$, but this contradicts P2.
So then, suppose that $x_1 \neq 0$. Then, we have

$$y \leq (1 - x_4) < (1 - x_3) < (1 - x_2) < (1 - x_1) \leq \max_{\theta \notin \text{supp}(\psi)} h(\theta), \tag{8}$$

i.e., $y < \max_{\theta \notin \text{supp}(\psi)} h(\theta)$, but this contradicts P2.
Finally, assume for contradiction that $\text{supp}(\psi) = [x_1, x_2]$, and consider $x_0 < x_1$ (i.e., $x_1 > 0$). We have $h(x_0) > 1 - x_1$, and then $h(x_0) > \min_{\theta \in \text{supp}(\psi)} h(\theta)$, contradicting P2.

Lemma 4. *If the buyer buys all ballots in* supp(ψ)*, then there is a canonical defense ψ' with the same value.*

Proof. Let ψ be a non-canonical defense. Suppose that $\text{supp}(\psi) \subseteq \mathcal{I}$, and let $d = \min_{\theta \in \text{supp}(\psi)} h(\theta)$. By Lemma 1, the buyer buys all ballots with $\theta \leq 1 - d$. Now let ψ' denote a canonical defense, and let $h'(\theta) = \frac{n_r f(\theta)(1-\theta)}{n_r f(\theta) + n_d \psi'(\theta)}$. Now $\min_{\theta \in \text{supp}(\psi')} h'(\theta) \geq d$ by P1. Thus, the buyer still buys all ballots with $\theta \leq 1-d$, including all of the decoys distributed according to ψ'.

Lemma 3 characterizes canonical defenses in terms of the properties defined above. Lemma 4 shows that if the buyer can buy up all decoys, then how they are distributed no longer matters.

Fixing the number of real ballots n_r, the EA's remaining choices are about n_d and ψ. We now state our main characterization result.

Theorem 1. *For a given n_r, n_d, and buyer budget B, the optimal strategy of the EA in the subgame perfect equilibrium is canonical.*

Proof. Assume for contradiction, that there is a non-canonical ψ that is better than any canonical defense. Let k be an index, and consider a sequence of defenses $\{\psi_k\} = \{\psi_0, \psi_1, ...\}$, where $\psi \stackrel{\text{def}}{=} \psi_0$. We will show that we can define a finite sequence that obtains a canonical defense at least as good as ψ. Let $h_k(\theta)$ denote the function h that corresponds to ψ_k.

Let $\mathcal{I}_k \subseteq [0,1]$ denote the set of intervals that are best for the buyer given ψ_k (solving for the buyer's objective subject to his budget). If the buyer buys all ballots in $\text{supp}(\psi_k)$, then by Lemma 4, we can modify ψ_k to form a canonical ψ_{k+1} with the same value, and we are done.

Suppose otherwise, and that in addition ψ_k does not satisfy P1 and P2. That is, we have:

(P0) the buyer does not buy all ballots in $\text{supp}(\psi_k)$, and one or both of
($\neg$ P1) $h_k(\theta)$ takes on multiple values for $\theta \in \text{supp}(\psi_k)$
($\neg$ P2) $\min_{\theta \in \text{supp}(\psi_k)} h_k(\theta) < \max_{\theta \notin \text{supp}(\psi_k)} h_k(\theta)$.

By P0, we can construct some interval $S_k \subseteq \text{supp}(\psi_k)$ (the *source set*), where the buyer is not buying all ballots, and an interval $T_k \subseteq \mathcal{I}_k$ (the *target set*), such that $h_k(S_k) < h_k(T_k)$ (and thus, $S_k \cap T_k = \varnothing$). Let $R_k = \text{supp}\psi \setminus \mathcal{I}_k$ be the remaining subset of $\text{supp}(\psi)$ that the buyer is not buying. We must have $\text{argmin}_{\theta \in \text{supp}(\psi_k)} h_k(\theta) \subseteq R_k$. The existence of T_k follows from $\neg P1$ because $\exists \theta \in \mathcal{I}_k$ for which $h_k(\theta) > \min_{\theta \in \text{supp}(\psi_k)} h_k(\theta)$ (the existence is guaranteed by values of $\theta \in \text{supp}(\psi_k)$ that are greater than the minimum), and thus we have $\max_{\theta \in \mathcal{I}_k} h_k(\theta) > \min_{\theta \in \text{supp}(\psi_k)} h_k(\theta)$. If $\neg P2$, then by buyer optimality (Lemma 1), $\text{argmin}_{\theta \in \text{supp}(\psi_k)} h_k(\theta) \subseteq R_k$. In both cases, $\text{argmin}_{\theta \in \text{supp}(\psi_k)} h_k(\theta) \subseteq S_k$.

We pick $\epsilon_S, \epsilon_T > 0$ to define a move of a uniform slice of ψ density from S_k to T_k such that,

(i) $\int_{\theta \in S_k} \max(0, \psi_k(\theta) - \epsilon_S)\, d\theta = \int_{\theta \in T_k} \epsilon_T\, d\theta$ [mass conservation]
(ii) $h_{k+1}(S_k) < h_{k+1}(T_k)$ [target set still preferred by buyer to source set]

By continuity (except possibly on a set of measure 0) of $h(\theta)$, such an ϵ_S, ϵ_T pair that satisfies (ii) exists. We argue that $S_k \cap \mathcal{I}_{k+1} = \varnothing$. Before the ψ mass is moved, we have $\min h_k(\mathcal{I}_k) \geq h_k(T_k) > h_k(S_k)$. After the move, we have $\min h_{k+1}(\mathcal{I}_{k+1}) \geq h_{k+1}(T_k) > h_{k+1}(S_k)$. The inequality is because the buyer can always exhaust his budget by buying $\mathcal{I}_k$. Thus, we know that the buyer does not buy anything in S_k after the ψ mass has been moved. Let $Q_k \stackrel{\text{def}}{=} \int_{\mathcal{I}_k} (1-\theta) f(\theta) d\theta$. Thus, we have $Q_{k+1} \leq Q_k$ because the only set on which $h_{k+1}(\theta) > h_k(\theta)$ is S_k. In addition, $\min_{\theta \in \text{supp}(\psi_k)} h_k(\theta) < \min_{\theta \in \text{supp}(\psi_{k+1})} h_{k+1}(\theta)$. Because $\forall k \in \mathbb{Z}^+$, $\theta \in [0,1]$, $h_k(\theta) \geq 0$ the sequence must be finite.

Note that $h_{k+1}(\theta)$ only differs from $h_k(\theta)$ at S_k and T_k, increasing at S_k and decreasing at T_k. We have

$$\begin{aligned}
&\min_{\theta \in \mathrm{supp}(\psi_{k+1})} h_{k+1}(\theta) \\
&= \min \left[\min_{\theta \in S_k} h_{k+1}(\theta), \min_{\theta \in T_k} h_{k+1}(\theta), \min_{\theta \in \mathrm{supp}(\psi_{k+1}) \setminus \{T_k, S_k\}} h_{k+1}(\theta) \right] \\
&> \min \left[\min_{\theta \in S_k} h_k(\theta), \min_{\theta \in S_k} h_k(\theta), \min_{\theta \in S_k} h_k(\theta) \right] = \min_{\theta \in \mathrm{supp}(\psi_k)} h_k(\theta). \qquad (9)
\end{aligned}$$

Theorem 1 says that for a given n_r and n_d, the optimal design of ψ by the EA is canonical. The next result shows that ψ (and its support, which is $[0, x_{\mathrm{O}}]$, "o" for optimal) can be easily computed given any n_r and n_d.

Theorem 2. *For any given n_r and n_d, the optimal defense of the EA in the subgame perfect equilibrium is given by a decoy ballot distribution with density function*

$$\psi(\theta) = \begin{cases} \frac{n_r}{n_d} \frac{(x_{\mathrm{O}} - \theta) f(\theta)}{1 - x_{\mathrm{O}}} & \text{for } \theta \in [0, x_{\mathrm{O}}] \\ 0 & \text{for } \theta \in (x_{\mathrm{O}}, 1] \end{cases}, \qquad (10)$$

where the threshold x_{O} is determined by the following equation: $\frac{1}{1-x_{\mathrm{O}}} \int_0^{x_{\mathrm{O}}} F(\theta) d\theta = \frac{n_d}{n_r}$ and $F(\theta)$ is the CDF of f.

Proof. We suppose that n_r and n_d are fixed, and solve the expression $h(\theta) = c$ for $\psi(\theta)$, where $\theta \in [0, x_{\mathrm{O}}]$ and $c > 0$, which gives us

$$\psi(\theta) = \frac{n_r}{n_d} \left(\frac{(1-\theta) f(\theta)}{c} - f(\theta) \right). \qquad (11)$$

Now, we need $\psi(\theta)$ to be non-negative on its support, which gives us $c \leq 1 - \theta, \forall \theta \in [0, x_{\mathrm{O}}]$, which implies that $c \leq 1 - x_{\mathrm{O}}$. Further, we need

$$\int_0^{x_{\mathrm{O}}} \frac{n_r}{n_d} \left(\frac{(1-\theta) f(\theta)}{1 - x_{\mathrm{O}}} - f(\theta) \right) = 1, \qquad (12)$$

which implies that $\frac{1}{1-x_{\mathrm{O}}} \left(F(x_{\mathrm{O}}) - \int_0^{x_{\mathrm{O}}} \theta f(\theta) d\theta \right) - F(x_{\mathrm{O}}) = \frac{n_d}{n_r}$, and after integrating by parts and using the fact that $\theta \geq 0$, we obtain $\frac{1}{1-x_{\mathrm{O}}} \int_0^{x_{\mathrm{O}}} F(\theta) d\theta = \frac{n_d}{n_r}$. Also, plugging in $1 - x_{\mathrm{O}}$ for c, we have, $\forall \theta \in [0, x_{\mathrm{O}}]$,

$$\psi(\theta) = \frac{n_r}{n_d} \left(\frac{(1-\theta) f(\theta)}{1 - x_{\mathrm{O}}} - f(\theta) \right) = \frac{n_r}{n_d} \frac{(x_{\mathrm{O}} - \theta) f(\theta)}{1 - x_{\mathrm{O}}}, \qquad (13)$$

as desired.

With this expression, we can determine the power of increasing the number of decoys, n_d, for any voter type distribution f, buyer budget B, and number of real ballots n_r.

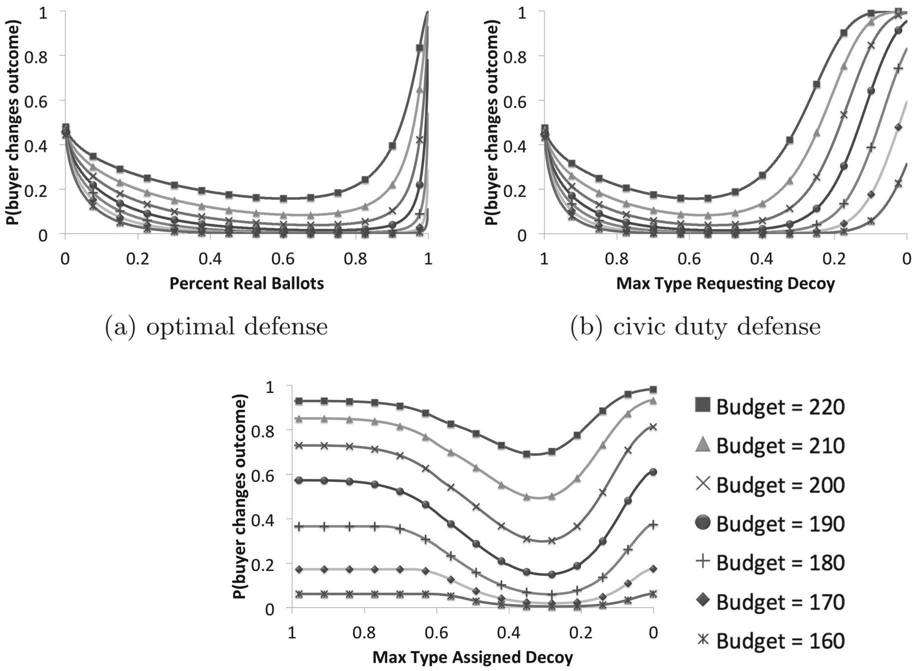

(a) optimal defense

(b) civic duty defense

(c) auction-based defense

Fig. 2. Comparing the power of different defenses, with f = Beta(2, 4), 1000 ballots in total (some real, some decoy), and different buyer budgets B. (a) Optimal defense, varying the fraction of real ballots. (b) Civic duty defense, with the EA optimizing the number of decoy ballots to use for each value of parameter x_{C} (the 'max type requesting decoy'). (c) Auction-based defense, with the EA optimizing the number of decoys to use for each value of x_{A} (the 'max type assigned a decoy').

5 Neutral Approaches

In this section, we consider defenses where the EA does not design ψ, since doing so may be argued as the EA playing too active a role in running the election. Beyond neutrality, these new approaches have the additional advantage of not relying on the EA having knowledge of f.

5.1 A Constrained Defense

We first consider a *constrained defense*:

Definition 2. *Defense ψ is* constrained *if the EA distributes decoy ballots uniformly at random, i.e., $\psi = f$.*

Having a constrained defense implies that $h(\theta) = \frac{n_r}{n_r+n_d}(1-\theta)$ and $\mathcal{I} = [0, \tau_C]$ for some $\tau_C > 0$, such that the budget is spent, i.e., $F(\tau_C) = B/(n_r + n_d)$.

Definition 3 (Low Budget). *A* low budget *is a budget where* $\int_{\tau_C}^{1} \theta f(\theta) d\theta < \frac{1}{2} - F(\tau_C)$.

Definition 4 (High Budget). *A* high budget *is a budget where* $\int_{\tau_C}^{1} \theta f(\theta) d\theta > \frac{1}{2} - F(\tau_C)$.

In words, for a buyer with a low (high) budget, the expected number of real ballots the buyer buys is lower than (exceeds) the amount needed to change the election outcome.

One way to study the power of a constrained defense is to consider the following question: if the total number of ballots is fixed, what is the optimal mix of real and decoy ballots?

Theorem 3. *Fixing the total number of ballots, the best constrained defense for the EA in the subgame perfect equilibrium is all (one) real ballots for low (high) buyer budget under the Normal approximation* (3).

Proof. We want to find, for fixed $n_r + n_d$,

$$\underset{\{n_r, n_d\}}{\operatorname{argmin}} P(\textit{buyer changes outcome}) \tag{14}$$

$$\approx \underset{\{n_r, n_d\}}{\operatorname{argmin}} P\left(\frac{\sqrt{n_r}(1 - 2F(\tau) - 2(1 - F(\tau))\mu_Y)}{2\sqrt{(1 - F(\tau))\mu_Y(1 - \mu_Y)}} < Z\right). \tag{15}$$

If a buyer has low budget, then this means that $\mu_Y(1 - F(\tau)) < \frac{1}{2} - F(\tau)$, which implies that

$$\frac{\sqrt{n_r}(1 - 2F(\tau) - 2(1 - F(\tau))\mu_Y)}{2\sqrt{(1 - F(\tau))\mu_Y(1 - \mu_Y)}} < 0, \tag{16}$$

and $P(\textit{buyer changes outcome})$ is minimized when $n_d = 0$. Similarly, if a buyer has high budget, then this means that $\mu_Y(1 - F(\tau)) > \frac{1}{2} - F(\tau)$, which implies that

$$\frac{\sqrt{n_r}(1 - 2F(\tau) - 2(1 - F(\tau))\mu_Y)}{2\sqrt{(1 - F(\tau))\mu_Y(1 - \mu_Y)}} > 0, \tag{17}$$

and $P(\textit{buyer changes outcome})$ is minimized when $n_r \to 0$.

With a low buyer budget, while a constrained defense makes the buyer buy some decoys, it also leaves unpurchased decoys and reduces the number of unpurchased real ballots, decreasing the accuracy of the result. Thus, decoys are not useful for the EA in this case. On the other hand, the best that the EA can do with a buyer with a high budget is to issue a single real ballot, with the hope that the buyer won't buy it, resulting in a high variance outcome based on the vote of a single voter. Decoys are used, but not to good effect.

5.2 Civic Duty Defense

In this model, the EA makes decoy ballots available to a random subset of those voters who make an explicit request for a decoy.[6] The decision of the EA is thus the number of decoy ballots, but not how to distribute them. Rather, this decision arises through a simple model of a societal process.

In modeling this process, we assume that, for a YES-buyer, there is some distribution of civic-mindedness $\pi(\theta)$, with support on $[0, x_C]$, that determines the probability that a voter will request a decoy, where x_C is a fixed, publicly known quantity ("*c*" for civic). In particular, we assume for simplicity that $\pi(\theta) \propto x_C - \theta$. This captures the idea that the more extreme an agent's type, the more likely the agent is to request a decoy and thus help preserve the election's integrity.

Via Bayes' rule, the effect on the distribution on types ψ of those who get decoys is $\psi(\theta) = P(\theta|\text{request decoy}) \propto P(\text{request decoy}|\theta)f(\theta) = \pi(\theta) \cdot f(\theta) = (x_C - \theta)f(\theta)$. In fact, there will sometimes be a choice of n_d such that the civic duty defense is optimal. If the EA can choose a number of decoys n_d such that $\frac{n_d(1-x_C)}{n_r} = k$, where k is the normalization constant, then we see the canonical structure, with $h(\theta) = 1 - x_C$, $\forall \theta \in [0, x_C]$. We call the defense obtained via this model a *civic duty defense*. An example of this defense is illustrated in Fig. 1(b).

5.3 Auction-Based Defense

In this variation, the EA makes decoy ballots available to voters via an auction. We assume a simple n_d+1st price auction (when selling n_d decoy ballots), with the EA choosing n_d. The intent is not to model a sophisticated auction, but to adopt a strategyproof mechanism as a model for an idealized market-based approach for distributing decoy ballots to voters. The effect is that decoys go to voters with the highest value for decoys. As with the civic duty defense, the EA who makes use of an auction-based defense chooses the number of decoy ballots but not how to distribute them.

In modeling this societal process, we assume that the value to a voter for a decoy is monotonically increasing as the voter's type θ gets closer to zero.[7] For this reason, we model the effect of the auction as being that there is some threshold $x_A \in (0, 1)$, whereby the decoys are distributed according to voter type distribution f, conditioned on $\theta \leq x_A$ ("A" for auction). In particular, for $\theta \in [0, x_A]$, we have $\psi(\theta) \propto f(\theta)$.

[6] We leave unmodeled that the buyer could try to interfere with this process. But notice that buying decoys from citizens who participate in this process is not useful because it depletes budget without hope of gaining real ballots. The same argument holds for the auction-based defense.

[7] We continue to assume that a voter's value for using a decoy is less than her value for a real ballot. Because of this, the auction-based process is consistent with our analysis in Sect. 2 in regard to the ordering of minimum acceptable offer price across different kinds of voters.

6 Simulation Results

We describe the results of an extensive simulation study to compare the power of various defenses in preventing a buyer succeeding in changing the outcome of an election. We choose to present results for voter type distribution $f = \text{Beta}(2, 4)$, but the analysis is qualitatively unchanged for other distributions, including those with mean voting types in $[0.01, 0.49]$.

Figure 4 fixes the number of real ballots, and shows that vote buying can be successfully thwarted by issuing sufficiently many decoy ballots. The optimal and civic duty defenses are most effective, but even issuing decoys according to the auction-based and constrained defenses substantially reduces the probability of a vote buyer's success. It is interesting that even a small number of decoys, relative to the number of real ballots, can be effective. It also helps with understanding to compare the power of different defenses when fixing the total number of ballots and varying the number of decoy ballots. Figure 2(a) shows the effect of varying the fraction of real ballots when using an optimal defense. Figures 2(b) and (c) show the effect of the civic duty defense and auction-based defence for different values of model parameter x_{C} (the 'max type requesting a decoy') and x_{A} (the 'max type winning a decoy'), with the EA optimizing the number of decoys for each value of x_{C} and x_{A}, respectively. The auction-based defense is the least effective, but even here there is a range of x_{A} for which the performance is better than without using any decoys. In Figs. 2(b) and (c), a maximum type of 0 receiving a decoy corresponds to zero decoys.

Fixing the total number of ballots, we can also examine the relative power of the different defenses as a function of the buyer budget. In Fig. 3 (with 1000 total ballots) we see that an optimal defense can use decoys to protect against buyers with around twice the budget of a 'no defense' approach that just uses real ballots. For the civic-duty and auction-based defenses, we fix $x_{\text{C}} = x_{\text{A}} = 0.5$ and pick the best n_d at each point in the graph. The auction-based defense is better than no defense and the constrained defense. The civic-duty defense has good performance, about that of the optimal defense for many buyer budgets.

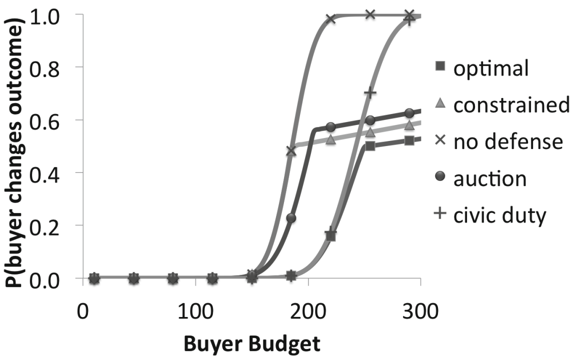

Fig. 3. Comparing the power of various defenses for $f = \text{Beta}(2, 4)$, x_{C} and $x_{\text{A}} = 0.5$, and 1000 total ballots.

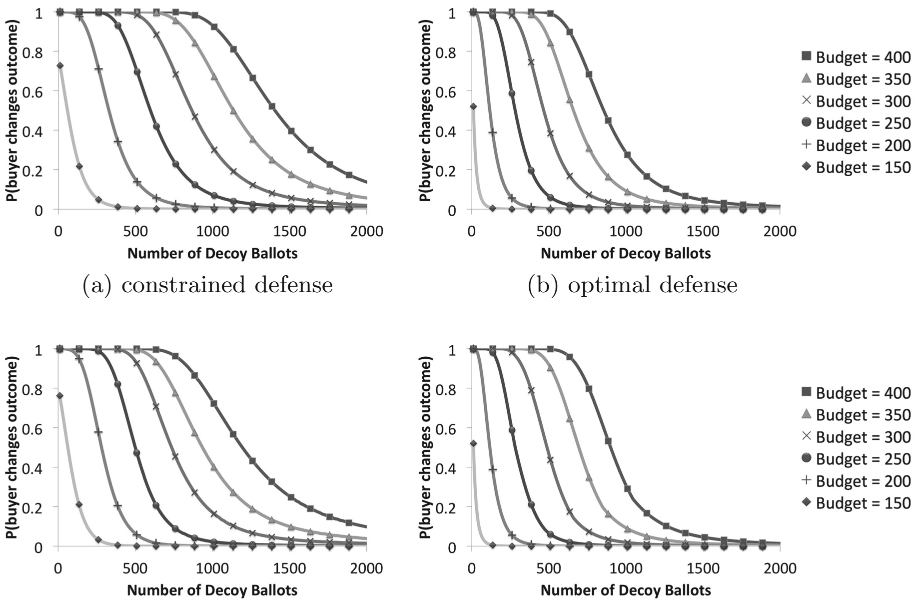

(a) constrained defense (b) optimal defense

(c) auction-based defense (d) civic duty defense

Fig. 4. Using decoys to thwart vote buying, for different buyer budgets (the number of ballots the buyer can buy). The number of real ballots is 750, the voter type distribution is f = Beta(2,4). (a) Constrained defense, in which decoy ballots are distributed according to $f(\theta)$. (b) Optimal defense. (c) Auction-based defense with $x_{\mathrm{A}} = 0.5$. (d) Civic duty defense with $x_{\mathrm{C}} = 0.5$.

7 Non-binary Election Outcomes

In this section, we consider a generalization of the model presented above to non-binary election outcomes. In particular, suppose that there are three election choices, X, Y, and Z, and assume, without loss of generality, that Z is expected to receive the most votes, followed by Y, followed by X. In this version, we consider the election outcome to be determined by plurality, although an alternative research direction could consider another rule such as single transferable vote.

There are three possible classes of buyers: an X-buyer, who wants the election outcome to be X, a Y-buyer, who wants the election outcome to be Y, and an XY-buyer, who wants the election outcome to be either X or Y. Here, we will discuss only X-buyers, leaving an analysis of the other two classes of buyers to future work.

We model voter types as being a vector of length 3, namely

$$\theta_i \stackrel{\text{def}}{=} (P(\text{voter } i \text{ votes } X), P(\text{voter } i \text{ votes } Y), P(\text{voter } i \text{ votes } Z)),$$

and use the shorthand $\theta_i[X]$, $\theta_i[Y]$, and $\theta_i[Z]$ to refer to the components of θ_i. Types are drawn from a distribution $g(\theta)$ with full support on a 2-simplex (e.g., a Dirichlet distribution or a discrete distribution with point masses).

Let $\mathcal{J} \subset [0,1] \times [0,1]$ denote the subset of types that the buyer buys. Let $\mathfrak{M} \stackrel{\text{def}}{=} \int_{\mathcal{J}} g(\theta) d\theta$ denote the fraction of real ballots that the buyer buys. Let the type distribution of the unbought types be given by

$$g_{\mathcal{T}}(\theta) \stackrel{\text{def}}{=} \begin{cases} \frac{g(\theta)}{1-\mathfrak{M}} & \text{for } \theta \in [0,1] \times [0,1] - \mathcal{J} \\ 0 & \text{for } \theta \in \mathcal{J} \end{cases}. \tag{18}$$

We use the notation $X_i \multimap f(\theta)$ to denote the hierarchical model $\theta_i \sim g(\theta); X_i \sim \text{Categorical}(\theta_i)$, and can now specify what it means for the buyer to change the election outcome. As in the proof of Lemma 2, let $V_i \multimap g_{\mathcal{T}}(\theta)$ denote the unbought votes, and let $W_j \multimap g(\theta)$ denote all votes. We then have

$$\begin{aligned} & P(\text{buyer changes outcome}) \\ &= P([\text{election bought}] \wedge [\text{correct outcome is } Y \text{ or } Z]) \\ &= P\left(\left[\frac{\sum_{i=1}^{(1-\mathfrak{M})n_r} \mathbb{1}_{V_i=X}}{n_r} + \mathfrak{M} > \max\left(\frac{\sum_{i=1}^{(1-\mathfrak{M})n_r} \mathbb{1}_{V_i=Y}}{n_r}, \frac{\sum_{i=1}^{(1-\mathfrak{M})n_r} \mathbb{1}_{V_i=Z}}{n_r}\right)\right]\right. \\ & \left.\wedge \left[\min\left(\frac{\sum_{j=1}^{n_r+n_d} \mathbb{1}_{W_j=Y}}{n_r+n_d}, \frac{\sum_{j=1}^{n_r+n_d} \mathbb{1}_{W_j=Z}}{n_r+n_d}\right) > \frac{\sum_{j=1}^{n_r+n_d} \mathbb{1}_{W_j=X}}{n_r+n_d}\right]\right), \end{aligned} \tag{19}$$

Recall that in the binary outcome case, we derived a simple characterization for an optimal buyer strategy (Lemma 1). On this basis, we were able to characterize the form of an optimal defense. In the three-outcome case, the strategy space is much richer, so we will discuss a few examples to illustrate some possible buyer strategies.

We first describe a simple vote buying strategy, and then show that it is optimal for simple, *deterministic types* (types where the voters vote for a particular outcome with probability 1).

Definition 5 (Buy the Expected Winner (BEW)). *The* buy the expected winner (BEW) *strategy is to greedily buy the type with the highest probability of voting for the current expected winner of the election, with the current expected winner determined considering the ballots already purchased by the buyer.*

Example 1. Suppose we have two voter types: 1,000 voters of type α: (0.25, 0, 0.75) and 600 voters of type β: (0, 1, 0) and no decoy ballots. Thus, the expected vote count is 250 X votes, 600 Y votes, and 750 Z votes. Table 1 illustrates the expected outcome for different buyer budgets and strategies. The third strategy for each budget above is to buy 200 votes of type α, and then buy four α votes for every three β votes until the budget runs out. This is the BEW strategy, which we can determine is optimal by enumerating all possible buyer strategies.

Example 2: Counterexample to Optimality of BEW. Suppose we have two voter types: 1,000 voters of type α: (0.25, 0, 0.75) and 1,000 voters of type β: (0, 0.26, 0.74) and no decoy ballots. The expected vote count is 250 X votes, 260 Y votes, and 1,490 Z votes. Table 2 illustrates the expected outcome for different buyer budgets and strategies. The buyer is better off buying type β than type α, which shows that the BEW strategy (i.e., buying type α) is not optimal. In fact, buying all type β is optimal, which can be seen by enumerating all possible buyer strategies.We next demonstrate that a refinement of BEW, where the buyer instead buys the type with highest $\max(\theta_i[Y], \theta_i[Z]) - \theta_i[X]$, can also be suboptimal.

Table 1. Illustrative buyer strategies for voter types α: (0.25, 0, 0.75) and β: (0, 1, 0).

Budget	Strategy	$E[\#X]$	$E[\#Y]$	$E[\#Z]$
0	-	250	600	750
400	Buy all α	550	600	450
400	Buy all β	650	200	750
400	Buy 314 α and 86 β	572.25	514	513.75
450	Buy all α	587.5	600	412.5
450	Buy all β	700	150	750
450	Buy 343 α and 107 β	614.25	493	492.75
500	Buy all α	625	600	375
500	Buy all β	750	100	750
500	Buy 371 α and 129 β	657.25	471	471.75

Table 2. Illustrative buyer strategies for voter types α: (0.25, 0, 0.75) and β: (0, 0.26, 0.74).

Budget	Strategy	$E[\#X]$	$E[\#Y]$	$E[\#Z]$
0	-	250	260	1,490
750	Buy all α	812.5	260	927.5
750	Buy all β	1,000	65	935

Example 3: Counterexample to Optimality of Refined BEW. Suppose we have 200 voters of type α: (0.25, 0.75, 0), 100 voters of type β: (0, 0.4, 0.6), and 150 voters of type γ: (0, 0, 1) and no decoy ballots. Then $E(\#X) = 50$, $E(\#Y) = 190$, $E(\#Z) = 210$. Suppose the buyer budget is 111. Table 3 illustrates the expected outcome for different buyer budgets and strategies. The first strategy

is the refined version of BEW. The buyer first buys 20 votes of type γ. Then, we have $E(\#X) = 70$, $E(\#Y) = 190$, $E(\#Z) = 190$. Now, he will buy 39 more ballots of type γ and 52 ballots of type α, resulting in $E(\#X) = 148$, $E(\#Y) = 151$, $E(\#Z) = 151$. So Y and Z are tied, and X has lost. The second strategy is to buy all 100 β votes and then 11 more γ votes. Here, we have $E(\#X) = 161$, $E(\#Y) = 150$, $E(\#Z) = 139$, and X wins. The third strategy, obtained by enumerating all possible strategies, is optimal.

Table 3. Illustrative buyer strategies for voter types α: (0.25, 0.75, 0), β: (0, 0.4, 0.6), and γ: (0, 0, 1).

Budget	Strategy	$E[\#X]$	$E[\#Y]$	$E[\#Z]$
0	-	50	190	210
111	Buy 52 α and 59 γ	148	151	151
111	Buy 100 β and 11 γ	161	150	139
111	Buy 2 α and 109 β	160.5	144.9	144.6

Example 4: Simple (Deterministic) Types. Suppose that we 200 voters of type X: (1, 0, 0), 350 voters of type Y: (0, 1, 0), and 450 voters of type Z: (0, 0, 1), and no decoy ballots. Table 4 illustrates the expected outcome for different buyer budgets and strategies. The third strategy is the BEW strategy, which is optimal here.

Table 4. Illustrative buyer strategies for deterministic voter types, X: (1, 0, 0), Y: (0, 1, 0), and Z: (0, 0, 1).

Budget	Strategy	$E[\#X]$	$E[\#Y]$	$E[\#Z]$
0	-	200	350	450
150	buy all Y	350	200	450
150	buy all Z	350	350	300
150	buy 25 Y and 125 Z	350	325	325

Example 5: Simple (Deterministic) Types with Decoys Suppose that the voter types are the same as in Example 4, but that the EA can issue decoys. We can numerically calculate the optimal EA strategy given the optimal defense. In regard to the optimal defense, this is BEW for some buyer budgets and numbers of decoys, but not always. See Fig. 5 for an illustration of the results. The optimal EA defense is to add the first 450 decoys with only Z type, and

then to begin adding both type Y and type Z decoys. In Fig. 5(a), all decoys are issued with type Z, and the optimal buyer strategy is to buy more Z ballots as each of them becomes less valuable— the buyer is playing the BEW strategy, now incorporating the probability that the ballots are real. In Fig. 5(b), some of the decoys (for numbers of decoys > 450) are issued with type Y, and the optimal buyer strategy is sometimes to buy all or nearly all Y ballots instead of Z ballots. In both cases, we see that the decoy defense is effective in stopping a vote buyer. The red line corresponds to the threshold where the buyer goes from winning in expectation to losing in expectation. With no defense, a strategic buyer needs a budget of 134 ballots to change the outcome of the election (where he would buy 17 Y ballots and 117 Z ballots). By issuing decoys, the EA can thwart a vote buyer with budgets including 150 and 300 (with 1,000 total real ballots).

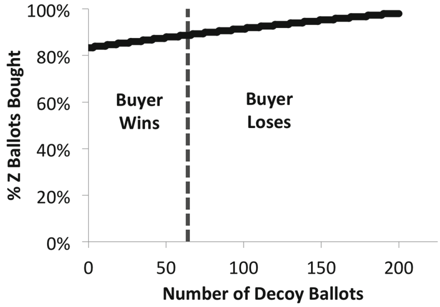

(a) buyer budget = 150 ballots

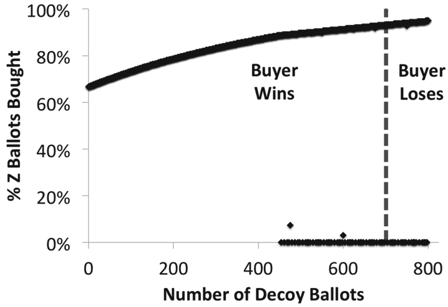

(b) buyer budget = 300 ballots

Fig. 5. Using optimally-distributed decoys to thwart vote buying in the three-outcome case. The voter types are from Example 5. In (a), the buyer is playing the BEW strategy, which is optimal. However, BEW is not optimal for all buyer budgets and numbers of decoys, as can be seen in (b), where the buyer sometimes buys all or nearly all Y ballots.

We can prove the optimality of BEW for these simple, deterministic types, and without decoy ballots. Note that with deterministic types there is no uncertainty about the outcome of the election.

Theorem 4. *For deterministic types and no decoy ballots, the BEW strategy is optimal for a buyer.*

Proof. We provide a proof for the slightly simpler case of buying fractional ballots (the proof for indivisible ballots follows the same outline). Let X, Y, and Z refer to the types (1, 0, 0), (0, 1, 0), and (0, 0, 1). Let x, y, and z refer to the number of ballots cast for each election outcome. We proceed to show that the BEW strategy minimizes the number of ballots needed for an X-buyer to change the outcome of the election. Let $\delta_y \geq 0, \delta_z \geq 0$ denote the number of Y and Z ballots purchased, respectively. The buyer wants to find the minimum $\delta = \delta_y + \delta_z$ s.t. $x + \delta_y + \delta_z \geq \max(y - \delta_y, z - \delta_z)$.

(Case 1) $x < y = z$. In BEW, the buyer buys Y and Z ballots in equal quantity until winning. In particular, buying $\delta_y^* = \delta_z^* = 1/3(y-x)$ leads to a win for X, since $x + \delta_y^* + \delta_z^* = x/3 + (2/3)y = y - \delta_y^* = z - \delta_z^*$. No strategy using $\delta' < \delta^* = \delta_y^* + \delta_z^* = 2/3(y-x)$ ballots can do better. We have

$$\min_{\delta_y', \delta_z' : \delta_y' + \delta_z' = \delta'} \max(y - \delta_y', z - \delta_z') \geq \min_{\delta_y, \delta_z : \delta_y + \delta_z = \delta^*} \max(y - \delta_y, z - \delta_z)$$
$$= \max(y - \delta_y^*, z - \delta_z^*) = x/3 + (2/3)y > x + \delta',$$

where the first inequality follows because the LHS is more constrained, and the first equality follows because this balances the two components of $\max(\cdot, \cdot)$.

(Case 2) $x < y < z$, and $1/2(x+z) \geq y$. In BEW, the buyer buys $\delta_y^* = 0$ and $\delta_z^* = 1/2(z-x)$ of the Y and Z ballots respectively. This leads to a win for X, with $x + \delta_z^* = z - \delta_z^* = (1/2)(z+x)$ (and $x + \delta_z^* = (1/2)(z+x) \geq y = y - \delta^*$.) No strategy using $\delta' < \delta^* = \delta_y^* + \delta_z^* = 1/2(z-x)$ ballots can do better. We have

$$\min_{\delta_y', \delta_z' : \delta_y' + \delta_z' = \delta'} \max(y - \delta_y', z - \delta_z') \geq \min_{\delta_y, \delta_z : \delta_y + \delta_z = \delta^*} \max(y - \delta_y, z - \delta_z)$$
$$= \max(y - \delta_y^*, z - \delta_z^*) = 1/2(x+z) > x + \delta',$$

where the first inequal. follows because the LHS is more constrained, and the first equality follows because $z - \delta^* \geq y$ and thus it is optimal to only buy Z ballots.

(Case 3) $x < y < z$, and $1/2(x+z) < y$ In BEW, the buyer first buys $z - y$ of the Z ballots, and then splits the remaining purchases equally between Y and Z ballots. In particular, $\delta_y^* = 1/3(y - (x + (z-y))) = 1/3(2y - x - z)$ and $\delta_z^* = (z-y) + 1/3(2y - x - z) = z - (y - \delta_y^*)$. Let $\delta^* = \delta_y^* + \delta_z^* = (1/3)y - (2/3)x + (1/3)z$. This leads to a win for X, with $x + \delta^* = (1/3)(x+y+z) = y - \delta_y^* = z - \delta_z^*$. No strategy using $\delta' < \delta^*$ ballots can do better. We have

$$\min_{\delta_y', \delta_z' : \delta_y' + \delta_z' = \delta'} \max(y - \delta_y', z - \delta_z') \geq \min_{\delta_y, \delta_z : \delta_y + \delta_z = \delta^*} \max(y - \delta_y, z - \delta_z)$$
$$= \max(y - \delta_y^*, z - \delta_z^*) = (1/3)(x+y+z) > x + \delta',$$

where the first inequality follows because the LHS is more constrained, and the first equality follows because this balances the two components of $\max(\cdot, \cdot)$.

An immediate corollary (noting that BEW is oblivious to budget) is that BEW also maximizes the advantage for X over the closest other outcome for a buyer with additional budget. We leave to future work to develop a full characterization of the optimal buyer strategy, and, in turn, optimal defense by the EA in the case of three or more outcomes. We do not yet have a characterization of the optimal buyer strategy even for the case of deterministic ballots, once decoys are also introduced.

8 Conclusion

We have presented the first game-theoretic study of the power of decoy ballots in thwarting vote buyers. We have characterized the form of an optimal defense, and compared its power to those of neutral defenses that could be enabled through leveraging simple societal processes to distribute decoy ballots. Our results are positive: decoy ballots are effective in thwarting the power of a vote buyer. Amongst the neutral defenses, the civic duty defense, where decoys are given at random to a subset of those who request such a ballot, seems especially interesting. Topics for future study include understanding defenses under the requirement that they must protect equally against a YES- or NO-buyer, and in settings with multiple buyers, simultaneous polls, and participants with value and cost heterogeneity. For the non-binary outcome case, we have provided some illustrative examples of the new subtleties that arise in modeling the optimal buyer strategy and thus optimal EA defense. There are a number of future directions of interest, including characterizing the optimal buyer and decoy defense strategies for non-binary outcome elections (initially for deterministic types). We expect that the richness of this setting will yield future interesting insights.

Acknowledgments. The Random Sample Voting (RSV) project has done much to encourage and facilitate this work, including convening a number of conference calls and circulating various internal documents. We are grateful to participants at EXPLORE 2017 for helpful discussions, as well as to the reviewers for their suggestions. LX acknowledges the support of the National Science Foundation, under grant IIS-1453542.

References

1. Binkele-Raible, D., Erdélyi, G., Fernau, H., Goldsmith, J., Mattei, N., Rothe, J.: The complexity of probabilistic lobbying. J. Discret. Optim. **11**, 1–21 (2014)
2. Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D. (eds.): Handbook of Computational Social Choice. Cambridge University Press, Cambridge (2016)
3. Bredereck, R., Chen, J., Hartung, S., Kratsch, S., Niedermeier, R., Suchý, O., Woeginger, G.J.: A multivariate complexity analysis of lobbying in multiple referenda. J. Artif. Intell. Res. **50**, 409–446 (2014)
4. Chaum, D.: Random-sample voting (2016). http://rsvoting.org/whitepaper/white_paper.pdf
5. Christian, R., Fellows, M., Rosamond, F., Slinko, A.: On complexity of lobbying in multiple referenda. Rev. Econ. Des. **11**(3), 217–224 (2007)
6. Dekel, E., Jackson, M.O., Wolinsky, A.: Vote buying: general elections. J. Polit. Econ. **116**(2), 351–381 (2008)
7. Faliszewski, P., Rothe, J.: Control and bribery in voting. In: Brandt, F., Conitzer, V., Endriss, A., Lang, J., Procaccia, A.D. (eds.) Handbook of Computational Social Choice. Cambridge University Press, Cambridge (2016)
8. Finan, F., Schechter, L.: Vote-buying and reciprocity. Econometrica **80**(2), 863–881 (2012)

9. Goel, A., Lee, D.T.: Towards large-scale deliberative decision-making: Small groups and the importance of triads. In: Proceedings of the 2016 ACM Conference on Economics and Computation, EC 2016, Maastricht, The Netherlands, pp. 287–303, 24–28 July 2016
10. Groseclose, T., Snyder, J.M.: Buying supermajorities. Am. Polit. Sci. Rev. **90**(2), 303–315 (1996)
11. Lee, D.T., Goel, A., Aitamurto, T., Landemore, H.: Crowdsourcing for participatory democracies: Efficient elicitation of social choice functions. In: Proceedings of the Second AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2014, 2–4 November 2014
12. Parkes, D.C., Wellman, M.P.: Economic reasoning and artificial intelligence. Science **349**(6245), 267–272 (2015)
13. Phongpaichit, P., Treerat, N., Chaiyapong, Y., Baker, C.: Corruption in the public sector in Thailand perceptions and experience of households. Political Economy Center. Chulalongkorn University, Bangkok (2000)
14. Schaffer, F.C. (ed.): Elections for Sale: The Causes And Consequences of Vote Buying. Lynne Rienner Publishers, London (2007)
15. Schaffer, F.C.: Why study vote buying? In: Schaffer, F.C. (ed.) Elections for Sale: The Causes And Consequences of Vote Buying, pp. 1–16. Lynne Rienner Publishers, London (2007)
16. Stokes, S.C., Dunning, T., Nazareno, M., Brusco, V.: Brokers, Voters, and Clientelism, chap. 8. What killed vote buying in Britain and the United States? Cambridge University Press, New York (2013)
17. Tambe, M.: Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned. Cambridge University Press, Cambridge (2011)
18. Vicente, P.C.: A model of vote-buying with an incumbency advantage (2013). http://www.pedrovicente.org/vb.pdf
19. Vicente, P.C.: Is vote buying effective? evidence from a field experiment in West Africa. Econ. J. **124**(574), F356–F387 (2014)