Thwarting Vote Buying Through Decoy Ballots

Citation

Parkes, David C., Paul Tylkin, and Lirong Xia. 2017. Thwarting Vote Buying Through Decoy Ballots. In AAMAS: Autonomous Agents and Multiagent Systems, São Paulo, Brazil, May 8-12, 2017, 45-66.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:41288103>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Thwarting Vote Buying Through Decoy Ballots

David C. Parkes

Harvard University
parkes@eecs.harvard.edu

Paul Tylkin

Harvard University
ptylkin@g.harvard.edu

Lirong Xia

Rensselaer Polytechnic Institute
xial@cs.rpi.edu

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.

1 Introduction

The goal of participatory democracy [Goel and Lee, 2016; Lee *et al.*, 2014] 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 “*one third of households were offered money to buy votes at the last general election*” [Phongpaichit *et al.*, 2000]. Schaffer [2007a] 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. He has also introduced the key notion of *proof of decoy* (see Section 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 200M 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 ballot cast is not ultimately counted, there is less emphasis on a voter needing to research an issue to form an opinion.

Although we situate our discussion in a societal context, similar themes can be imagined for economies of AIs [Parkes

¹In some approaches, this cost comes about, in addition, as a result of needing to physically mail ballots [Chaum, 2016].

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

Figure 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 $f = \text{Beta}(1, 2)$.

and Wellman, 2015], 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.

1.1 Our Contributions

We provide a formal model of vote buying, including a characterization of the vote buyer’s 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.

1.2 Related Work

There are numerous studies on vote buying, for example [Stokes *et al.*, 2013; Vicente, 2014; Finan and Schechter, 2012; Schaffer, 2007b]. These include game-theoretic models of vote buying, but none that consider the role of decoy ballots. In Dekel *et al.* [2008], the game is played by the candidates themselves buying votes, Groseclose and Snyder [1996] study vote buying in legislative bodies and analyze the optimal coalition size. Vicente [2013] 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 [Brandt *et al.*, 2016; Faliszewski and Rothe, 2016]. 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 [Christian *et al.*, 2007; Brederick *et al.*, 2014; Binkele-Raible *et al.*, 2014]. 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* [Tambe, 2011], 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 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

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

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

Figure 2: Comparing the power of different defenses, with $f = \text{Beta}(2, 4)$, 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’).

“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.

(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 prove to anyone that she has

a decoy. 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 [2016].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:

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

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. 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

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

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 .

⁴To understand this ordering, consider 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 change their vote and vote YES. Thus, these are the most expensive votes to buy. Analogously, 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 higher value for depleting the budget of a YES-buyer than a voter who would vote YES.

⁵The asymmetry in proof-of-decoy but no proof-of-authenticity is important in preventing a buyer from using coercion to buy only real ballots, while at the same time allowing a voter with a decoy ballot to sell with impunity to accusations of acting against the social good (since she can, if challenged to do so, prove that it is decoy, and thus that she is acting in good faith.) A voter will never choose to reveal that she holds a decoy to a buyer, since doing so would just cause the buyer to refuse to transact with her.

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.*

Where proofs are omitted, this is because of space.⁶

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 $P(\text{buyer changes outcome})$*

$$= P([\text{election bought}] \wedge [\text{correct outcome is NO}]) \approx P\left(\frac{n_r(1 - 2\mathfrak{M} - 2(1 - \mathfrak{M})\mu_Y)}{2\sqrt{n_r(1 - \mathfrak{M})\mu_Y(1 - \mu_Y)}} < Z < \frac{(1 - 2\mu)\sqrt{n_r + n_d}}{2\sqrt{\mu(1 - \mu)}}\right), \quad (2)$$

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

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.

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 Figure 1(a) for an illustration of a canonical defense. Let $\text{supp}(g)$ 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.*

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

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.*

⁶The longer version of the paper, which includes all proofs, is available at request from the authors.

Figure 3: Comparing the power of various defenses for $f = \text{Beta}(2, 4)$, x_c and $x_A = 0.5$, and 1000 total ballots.

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 $\{\psi_k\} = \{\psi_0, \psi_1, \dots\}$, 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, $S_k \cap T_k = \emptyset$). 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 $\arg \min_{\theta \in \text{supp}(\psi_k)} h_k(\theta) \subseteq R_k$. 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), $\arg \min_{\theta \in \text{supp}(\psi_k)} h_k(\theta) \subseteq R_k$. In both cases, $\arg \min_{\theta \in \text{supp}(\psi_k)} h_k(\theta) \subseteq S_k$.

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 $S_k \cap \mathcal{I}_{k+1} = \emptyset$. 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. $\square$

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 $[0, x_o]$, “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*

$$\psi(\theta) = \begin{cases} \frac{n_r}{n_d} \frac{(x_o - \theta)f(\theta)}{1 - x_o} & \text{for } \theta \in [0, x_o] \\ 0 & \text{for } \theta \in (x_o, 1] \end{cases}, \quad (3)$$

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

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 .

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 (2).*

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.⁷ 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 Figure 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+1 st 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.⁸ For this reason, we model the

⁷We 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.

⁸We insist, though, that the reasonable property holds that a voter’s value for using a decoy is less than her value for a real ballot, and thus this auction-based societal process is consistent with our analysis in Section 2 in regard to the ordering of minimum acceptable offer price from a buyer across different kinds of voters.

Figure 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 $f = \text{Beta}(2,4)$. (a) Constrained defense, in which decoy ballots are distributed according to $f(\theta)$. (b) Optimal defense. (c) Auction-based defense with $x_A = 0.5$. (d) Civic duty defense with $x_C = 0.5$.

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 Simulation Results

We describe the results of an extensive simulation study to compare 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 2(c) show the effect of the civic duty defense and auction-based defense 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 Figures 2(b) and 2(c), a maximum type of 0 receiving a decoy corresponds to zero decoys. Also fixing the total number of ballots, we ex-

amine the relative power of the different defenses as a function of the buyer budget. In Figure 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 $x_C = x_A = 0.5$ 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, almost the same as the optimal defense for many buyer budgets.

7 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 more than two ballot choices, multiple buyers, simultaneous polls, and participants with value and cost heterogeneity.

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. LX acknowledges the support of the National Science Foundation, under grant IIS-1453542.

References

- [Binkele-Raible *et al.*, 2014] Daniel Binkele-Raible, Gábor Erdélyi, Henning Fernau, Judy Goldsmith, Nicholas Mattei, and Jörg Rothe. The Complexity of Probabilistic Lobbying. *Discret. Optim.*, 11, February 2014.
- [Brandt *et al.*, 2016] Felix Brandt, Vincent Conitzer, Ulle Endriss, Jerome Lang, and Ariel D. Procaccia, editors. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- [Bredereck *et al.*, 2014] Robert Bredereck, Jiehua Chen, Sepp Hartung, Stefan Kratsch, Rolf Niedermeier, Ondrej Suchý, and Gerhard J. Woeginger. A Multivariate Complexity Analysis of Lobbying in Multiple Referenda. *J. Artif. Intell. Res. (JAIR)*, 50:409–446, 2014.
- [Chaum, 2016] David Chaum. Random-Sample Voting. http://rsvoting.org/whitepaper/white_paper.pdf, 2016.
- [Christian *et al.*, 2007] Robin Christian, Mike Fellows, Frances Rosamond, and Arkadii Slinko. On complexity of lobbying in multiple referenda. *Review of Economic Design*, 11(3):217–224, 2007.
- [Dekel *et al.*, 2008] Eddie Dekel, Matthew O. Jackson, and Asher Wolinsky. Vote Buying: General Elections. *Journal of Political Economy*, 116(2):351–381, 2008.
- [Faliszewski and Rothe, 2016] Piotr Faliszewski and Jörg Rothe. Control and bribery in voting. In *Handbook of Computational Social Choice*, chapter 7. Cambridge University Press, 2016.
- [Finan and Schechter, 2012] Frederico Finan and Laura Schechter. Vote-Buying and Reciprocity. *Econometrica*, 80(2):863–881, 2012.
- [Goel and Lee, 2016] Ashish Goel and David Timothy Lee. 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 '16, Maastricht, The Netherlands, July 24-28, 2016*, pages 287–303, 2016.
- [Groseclose and Snyder, 1996] Tim Groseclose and James M. Snyder. Buying supermajorities. *The American Political Science Review*, 90(2):303–315, 1996.
- [Lee *et al.*, 2014] David Timothy Lee, Ashish Goel, Tanja Aitamurto, and Hélène Landemore. Crowdsourcing for participatory democracies: Efficient elicitation of social choice functions. In *Proceedings of the Second AAAI Conference on Human Computation and Crowdsourcing, HCOMP 2014, November 2-4, 2014, Pittsburgh, Pennsylvania, USA*, 2014.
- [Parkes and Wellman, 2015] David C. Parkes and Michael P. Wellman. Economic Reasoning and Artificial Intelligence. *Science*, 349(6245):267–272, 2015.
- [Phongpaichit *et al.*, 2000] Pasuk Phongpaichit, Nualnoi Treerat, Yongyuth Chaiyapong, and Chris Baker. Corruption in the Public Sector in Thailand Perceptions and Experience of Households. Political Economy Center (Bangkok: Chulalongkorn University), 2000.
- [Schaffer, 2007a] Frederic Charles Schaffer, editor. *Elections for Sale: The Causes And Consequences of Vote Buying*. Lynne Rienner Pub, 2007.
- [Schaffer, 2007b] Frederic Charles Schaffer. Why Study Vote Buying? In *Elections for Sale: The Causes And Consequences of Vote Buying*, chapter 1, pages 1–16. Lynne Rienner Pub, 2007.
- [Stokes *et al.*, 2013] Susan C. Stokes, Thad Dunning, Marcelo Nazareno, and Valeria Brusco. *Brokers, Voters, and Clientelism*, chapter 8. What Killed Vote Buying in Britain and the United States? Cambridge University Press, 2013.
- [Tambe, 2011] Milind Tambe. *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned*. Cambridge University Press, 2011.
- [Vicente, 2013] Pedro C. Vicente. A Model of Vote-buying with an Incumbency Advantage. <http://www.pedrovicente.org/vb.pdf>, 2013.
- [Vicente, 2014] Pedro C. Vicente. Is Vote Buying Effective? Evidence from a Field Experiment in West Africa. *The Economic Journal*, 124(574):F356–F387, 2014.