

Doodle Poll Games

Svetlana Obraztsova
Nanyang Technological University
Singapore
lana@ntu.edu.sg

Zinovi Rabinovich
Nanyang Technological University
Singapore
zinovi@ntu.edu.sg

Maria Polukarov
University of Southampton
Southampton, UK
mp3@ecs.soton.ac.uk

Edith Elkind
University of Oxford
Oxford, UK
elkind@cs.ox.ac.uk

ABSTRACT

In Doodle polls, each voter approves a subset of the available alternatives according to his preferences. While such polls can be captured by the standard models of Approval voting, Zou et al. [18] analyse real-life Doodle poll data and conclude that poll participants' behaviour seems to be affected by considerations other than their intrinsic preferences over the alternatives. To capture this phenomenon, they propose a model of social voting, where voters approve their top alternatives as well as additional 'safe' choices so as to appear cooperative. The predictions of this model turn out to be consistent with the real-life data. However, Zou et al. do not attempt to rationalise the voters' behaviour in the context of social voting: they explicitly describe the voters' strategies rather than explain how these strategies arise from voters' preferences. In this paper, we complement the work of Zou et al. by putting forward a model in which the behaviour described by Zou et al. arises as an equilibrium strategy. In our model, a voter derives a bonus from approving each additional alternative, up to a certain cap. We show that trembling hand perfect Nash equilibria of our model behave consistently with the model of Zou et al. Importantly, placing a cap on the total bonus is an essential component of our model: in the absence of the cap, all Nash equilibria are very far from the behaviour observed in Doodle polls.

Keywords

Trembling Hand Nash Equilibrium; Approval voting; Doodle polls

1. INTRODUCTION

Scheduling group meetings is a tedious and time-consuming activity: the participants have to decide which time slots are suitable for all or most of them, and to choose one or more slots based on these constraints. A convenient online tool for this task is *Doodle*: the poll initiator creates a list of possible time slots, and all participants can then cast their vote online, indicating which time slots are acceptable for them. In the most basic version of a Doodle poll, everyone can mark each time slot as suitable or not (there is also a more flexible option, where each participant is allowed three levels of approval: 'yes', 'no', and 'if need be'), and each participant can vote 'yes' for any number of slots.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)*, S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.
Copyright © 2017, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Assuming for simplicity that in the end the meeting time will be chosen among the time slots that received the largest number of votes (using a deterministic or randomised tie-breaking rule), we can view this situation as an approval voting scenario. Now, in case of dichotomous preferences (where each voter either approves or disapproves each time slot), approval voting is known to elicit truthful behaviour from the voters: it is a weakly dominant strategy for each voter to vote exactly for the time slots he approves. Of course, in reality voter preferences may be more complicated: for instance, even if one currently has no appointments scheduled for Monday 08:30, one may prefer not schedule a meeting for this time if at all possible. If each voter has a weak order over the time slots, a relevant notion is that of a *sincere* strategy: an approval vote is said to be sincere if the voter weakly prefers each of the approved alternatives to each of the non-approved alternatives (see, e.g., [4]). Under reasonable assumptions on tie-breaking (and, in case of randomised tie-breaking, voters' preferences over the respective lotteries), approval voting is known to encourage such sincere behaviour [7].

However, these theoretical results do not seem to fully explain the voters' behaviour that is observed in practice; in particular, they provide no clue as to how the voters decide how many alternatives to approve. To address this challenge, recently Zou et al. [18] used a large dataset from Doodle (over 14 million votes from 2 million participants in over 340,000 polls) to analyse user behaviour in Doodle polls. Their data shows that participants tend to behave differently in *closed polls* (where previously cast votes are not visible) and in *open polls* (where previously cast votes can be seen). Specifically, participants in open polls tend to approve more slots and coordinate with the previous voters. Perhaps more intriguingly, in open polls both the most popular slots and the least popular slots tend to receive more votes than in closed polls, with medium popularity slots receiving a similar number of votes.

Zou et al. explain this phenomenon by introducing the model of *social voting*. In this model, besides having intrinsic preferences for different time slots, voters would like to appear cooperative, and therefore they gain extra utility from approving many slots. Thus, in addition to approving their most preferred slots, they may approve a few extra slots. It makes sense for them to choose these slots among the unpopular slots (those that have received few votes so far), to minimise the risk that these slots will actually be selected. Zou et al. show that this model is consistent with the data: they generate synthetic data according to their model and obtain voter behaviour that is qualitatively similar to what is observed in Doodle open polls.

However, the social voting model proposed by Zou et al. makes no attempt to rationalise the voters' behaviour: it simply stipulates

that each voter will vote for his most preferred slots and some of his somewhat less preferred slots based on their popularity; the authors do not reconstruct voters’ utilities that make such behaviour *rational*. The goal of our work is to fill this gap.

In our model, each voter assigns some utility to each time slot (with different utilities corresponding to different *preference levels*), and the winning alternative is selected among the alternatives with the highest number of approval points using either a lexicographic tie-breaking rule (based on a fixed ordering of the alternatives) or the randomised tie-breaking rule (where the winner is chosen among the top-scoring alternatives uniformly at random). We remark that tie-breaking rules play a particularly important role in Doodle polls: since the number of alternatives is often large relative to the number of votes, ties are quite likely, and, since the decisions made by Doodle polls are typically low-stake, both tossing a fair coin and picking the lexicographically first among the top-scoring alternatives is socially acceptable. It remains to explain how voters derive utility from appearing cooperative.

A natural first attempt is to assume that, besides deriving the respective utility from the winning slot, each voter obtains a very small bonus from each time slot he approves; the second component is supposed to capture his social utility. However, we show that under these assumptions in each Nash equilibrium (NE) of the resulting game each candidate is approved by almost all voters, and, moreover, if there is no consensus among the voters, the existence of Nash equilibria requires a very delicate balance of voters’ preferences. Thus, pure Nash equilibria in this model cannot be viewed as realistic predictions of voters’ behaviour.

We then consider a more refined version of this approach, where a voter only gets the social bonus for the first κ time slots that he approves, for a value of κ that is noticeably smaller than the total number of alternatives. Indeed, it is plausible that, in practice, to appear cooperative, one only needs to approve a number of alternatives that exceeds a certain threshold, and approving further alternatives may not contribute much to the voter’s perception by his peers. This is corroborated by empirical findings, which demonstrate that voters seem to approve very popular and very unpopular, but not all available slots [18]. This model admits a much richer set of pure Nash equilibria; indeed, we obtain a plethora of ‘bad’ Nash equilibria where everyone approves one alternative (which may be viewed as undesirable by some or even all voters) and $\kappa - 1$ other alternatives, and approvals are distributed so that every non-winning alternative is far from becoming a winner; a similar phenomenon is well-known in the context of Plurality voting (see, e.g., [14]).

Therefore, it becomes important to come up with a suitable *equilibrium refinement* that will help us identify Nash equilibria that are more likely to occur in practice. To do so, we employ a variant of Selten’s celebrated *trembling hand perfect Nash equilibrium (THPE)* [16, 12]: under this notion, each player assumes that, with small probability, other players may make a mistake when implementing their strategy and choose another (random) strategy instead. It has been argued that THPE is the most important refinement of Nash equilibrium [17]; moreover, it is particularly appealing in our setting, where the number of decisions that a voter needs to make may be large, and mistakes are rarely costly. Trembling hand perfect equilibria have been studied for a wide variety of games; in particular, in the context of voting, different variants of THPE have been proposed to explain voter behaviour in plurality and runoff rule voting games [13, 15], information revelation scenarios [11] and agenda-setting games [1].

In this work, we modify Selten’s original definition by restricting the type of mistakes that players may make: specifically, we assume that each player makes a mistake independently for each

alternative, i.e., for each alternative he considers, with probability ε he votes ‘Yes’ if he intends to disapprove this alternative and ‘No’ if he intends to approve it, and he votes correctly with probability $1 - \varepsilon$. We then show that, in games with capped bonuses, a voter’s best response to a trembling hand strategy of other players is to approve all the alternatives at his top preference level as well as some of the most unpopular alternatives, with the total number of approved alternatives never exceeding the cap. This result is consistent with the findings of [18]. The analysis of this model is quite involved, and can be seen as our main technical contribution.

While our primary goal is to explain voters’ behaviour in Doodle polls, we expect the class of approval voting games with a social bonus to be of broader interest. Hence, when analysing such games, we do not limit ourselves to the regime that provides the closest approximation to Doodle poll games. In particular, while our primary focus is the setting where the social bonus is capped, we also analyse the setting with uncapped social bonus. For the latter setting we provide efficient algorithms for computing Nash equilibria if ties are broken lexicographically or the voters’ true preferences are dichotomous. We complement these results by showing that, for uncapped social bonus, it is NP-hard to decide if a given preference profile admits a pure strategy Nash equilibrium if ties are broken in a randomised fashion and voters have three or more preference levels. We hope that this analysis will prove useful when considering other applications of approval voting.

From a conceptual perspective, our contribution is twofold. First, while the notion of a social bonus was put forward by [18], we are able to show that this concept can be used to *rationalise* voters’ behaviour. This turns out to be a non-trivial task: the straightforward approach of incorporating the social bonus into the voters’ utilities, in an uncapped fashion, results in a model with very few Nash equilibria, which are, moreover, rather unnatural. To overcome this difficulty, we introduce the idea of a cap on the social bonus, which is both intuitively appealing and enables us to obtain results that agree with the real-life data. Another innovation is the use of trembling hand perfection as an equilibrium refinement tool: while this is an elegant and conceptually appealing notion, it has received surprisingly little attention in the algorithmic game theory literature; notable exceptions are papers by Hansen, Etesami et al. [10, 8], which, unlike our work, focus on abstract normal-form games and obtain computational hardness results rather than efficient algorithms, as well as a very recent paper on trembling hand equilibria of plurality voting [15], which differs from this work in that it considers a different voting rule (plurality rather than approval) and no social bonuses.

2. PRELIMINARIES

The model we introduce in this paper extends the standard model of approval voting, so we start by formally defining approval voting and other relevant concepts.

In approval voting, there is a set $V = \{v_1, v_2, \dots, v_{|V|}\}$ of voters, electing a single winner from a set $C = \{c_1, c_2, \dots, c_{|C|}\}$ of alternatives, or candidates. A single vote (or ballot) of voter $v \in V$ is a subset of candidates $b^v \subseteq C$ that he approves. We will also regard b^v as a $|C|$ -dimensional binary vector $(b_1^v, \dots, b_{|C|}^v)$ with $b_c^v = 1$ if v approves alternative $c \in C$ and $b_c^v = 0$ otherwise. A voting profile $\mathbf{b} = (b^v)_{v \in V}$ is a vector of ballots, one for each voter.

Given a profile \mathbf{b} , let $s_c(\mathbf{b}) = \sum_{v \in V} b_c^v$ denote the score of candidate c ; the vector $s(\mathbf{b}) = (s_c(\mathbf{b}))_{c \in C}$ is then the score vector of \mathbf{b} . The set of (provisional) winners of the election, or the election outcome, is given by the set of alternatives $W(\mathbf{b}) = \arg \max_{c \in C} s_c(\mathbf{b})$; the ties among the provisional winners are bro-

ken by a given *tie-breaking rule*. We consider two common tie-breaking rules: *lexicographic*, which selects the lexicographically first candidate in $W(\mathbf{b})$ with respect to a given linear order over C , and *randomised*, which chooses one uniformly from $W(\mathbf{b})$.

Each voter $v \in V$ has *preferences over individual candidates*, \geq_v , modelled as a total (but not necessarily strict) order on C . That is, \geq_v is a reflexive, transitive and complete (but not necessarily anti-symmetric) binary relation on C . We write $x \geq_v y$ to express that voter v likes candidate x at least as much as candidate y . We write $x >_v y$ (strict preference) if $x \geq_v y$ but not $y \geq_v x$. Thus, \geq_v defines a partition of C into L disjoint subsets $\{C_1^v, \dots, C_L^v\}$, so that v is indifferent among the alternatives in the same element of the partition, but strictly prefers any alternative in C_ℓ^v over any alternative in C_k^v , for all $\ell > k$. That is, L denotes the number of *preference levels* of the voter. We allow for the possibility that some, but no more than $L - 2$, elements of the partition are empty. That is, no voter is indifferent between all the alternatives in C , and it is without loss of generality to assume that all voters have the same number of preference levels L . Specifically, L is the maximum number of proper (i.e., non-empty) preference levels across all the voters, and we require that for each voter v the preference levels 1 and L are not empty. A voter with exactly two proper preference levels is called *dichotomous*, and one with three proper preference levels is called *trichotomous*.

When analysing voters' strategic behaviour, we need to reason about voters' preferences over election outcomes, i.e., over sets of provisional winners. Note that under deterministic tie-breaking, comparing every pair of outcomes is easy: we apply the tie-breaking rule to determine the eventual winner in each set, and compare these winners using the voter's preference order \geq_v . However, if ties are broken randomly, \geq_v does not induce a complete order over all possible outcomes, and a common solution (see, e.g., [5, 2, 9, 6, 15]) is to augment voters' preferences with cardinal utilities. To this end, we assume that each voter v is endowed with a function $\delta_v : C \rightarrow \mathbb{Q}$ that takes at most L distinct values and is consistent with v 's preferences: $\delta_v(c) > \delta_v(c')$ if and only if $c >_v c'$. Without loss of generality, we assume that $\delta_v(c) = 1$ if $c \geq_v c'$ for all $c' \in C$ and $\delta_v(c) = 0$ if $c' \geq_v c$ for all $c' \in C$. Now, the *utility from the outcome*, $u_v(\mathbf{b})$, of voter v under ballot profile \mathbf{b} is defined as follows. If the ties are broken deterministically, according to a strict linear order \succ , then

$$u_v(\mathbf{b}) = \delta_v(c_{\max}), \text{ where} \quad (1)$$

$$c_{\max} \in W(\mathbf{b}) \text{ and } c_{\max} \succ c \text{ for all } c \in W(\mathbf{b}) \setminus \{c_{\max}\}.$$

In case of uniform tie-breaking, we consider the expected utility:

$$u_v(\mathbf{b}) = \frac{1}{|W(\mathbf{b})|} \sum_{c \in W(\mathbf{b})} \delta_v(c). \quad (2)$$

Assuming that each voter strives to maximise his utility from the outcome, for a fixed tie-breaking rule our setting induces a normal-form game $\Gamma = \langle V, 2^C, (u_v)_{v \in V} \rangle$ where the set of players is given by the set of voters V , the strategy set available to each player is given by the collection of all subsets (i.e., the power set) of the set of candidates C , and the utility function of each player is given by u_v as defined by equations (1) or (2) above. We call Γ the *approval voting game* (the tie-breaking rule will be clear from the context).

Given voter v 's preference order \geq_v , a ballot b^v is called *sincere* if the voter prefers each approved candidate to each disapproved candidate; that is, if $x \geq_v y$ for all $x \in b^v$ and all $y \in C \setminus b^v$ [4]. Observe that according to this definition, a voter can vote sincerely in a number of ways, and abstention ballots, in particular, are also considered sincere. It is known that if a voter's preferences over in-

dividual alternatives are dichotomous, then he always has a sincere best response under approval voting, irrespective of how these individual preferences are extended to preferences over outcomes (i.e., sets)—that is, irrespective of how the utility function of game Γ is defined [3]. Under some additional assumptions on set preferences, this claim also holds for trichotomous individual preferences, but if a voter has four or more levels of preference, he may prefer to vote insincerely. Recently, Endriss [7] has formulated several principles of lifting individual preferences to set preferences under which a voter will always have a sincere best response in an approval voting game, even if he has more than three proper preference levels. Importantly, the utility functions defined by equations (1) and (2) satisfy these principles.

3. MODEL

Doodle polls can be seen as an implementation of the approval voting rule, which was discussed in the previous section. Specifically, candidates are time slots, and voters indicate their (un)availability at these slots. We note that Doodle also allows voters to express trichotomous preferences, by classifying the time slots into ones that are (1) convenient, (2) feasible but inconvenient or (3) not feasible; however, this setting is not captured by the standard model of approval voting, so we will not consider it in our work.

Our goal is to provide a game-theoretic model that explains the experimental findings of Zou et al. [18], and in particular the phenomenon of social voting. To this end, we incorporate social rewards into voters' utility functions. Specifically, we assume that voter v gets a *social bonus*, β , for each of his approved alternatives, as long as their number does not exceed the *cap* κ , $0 \leq \kappa \leq |C|$. This cap is the maximum number of approved candidates that is socially rewarded. In approval games, $\kappa = 0$; if $\kappa = |C|$, a voter gets a bonus for each alternative he approves. We assume that social bonuses never prevail over the original preferences:

$$0 < \beta < \frac{1}{|C|} \min_{v \in V} \min_{c, c' \in C} |\delta_v(c) - \delta_v(c')|. \quad (3)$$

The *total utility* of voter v under ballot profile \mathbf{b} is then composed of the utility from the outcome $u_v(\mathbf{b})$ (given by equation (1) or (2), as per the tie-breaking rule), and the overall social bonus:

$$U_v(\mathbf{b}) = u_v(\mathbf{b}) + \beta \cdot \min\{|b^v|, \kappa\}. \quad (4)$$

DEFINITION 1. A Doodle poll game (DPG) is a normal-form game $\Gamma = \langle V, 2^C, (U_v)_{v \in V} \rangle$ with the set of players V , where for each player v , his strategy set is the power set of the set of alternatives C and his utility function U_v is defined by equation (4).

4. UNCAPPED SOCIAL BONUS

We start by providing a few simple observations about the structure of equilibrium profiles in DPGs with uncapped social bonus. We focus on pure strategy Nash equilibria, to which we may refer simply as equilibria or Nash equilibria (NE). Our first observation applies irrespective of whether there is a cap on the social bonus.

OBSERVATION 1. For both of our tie-breaking rules, in each DPG every voter has a best response where he approves all the alternatives at his highest level of preference (and possibly some other alternatives).

Indeed, if a voter fails to approve some of his most preferred alternatives, he can only increase his total utility by adding such an alternative to his ballot: by doing so he cannot lower his utility from the outcome, while his social bonus may only grow (unless capped). In fact, for $\kappa = |C|$, every best response is of this form.

OBSERVATION 2. For both of our tie-breaking rules, in every equilibrium of a DPG with $\kappa = |C|$ every voter approves all the alternatives at his highest level of preference.

When $\kappa = |C|$ and ties are broken lexicographically, it is beneficial for everyone to vote for the election winner.

OBSERVATION 3. In any equilibrium of a DPG with $\kappa = |C|$ and lexicographic tie-breaking, all voters approve the (unique) winner of the election.

Indeed, if there is a voter that fails to approve the current winner, he will strictly increase his utility by adding the winner to his ballot: this move will not change the winner—and hence the utility from the outcome—for the voter, whereas his social bonus will grow. As for the other candidates, for similar reasons, in an equilibrium profile they all must get maximum possible support, but so that the current winner remains unchanged.

OBSERVATION 4. In any equilibrium of a DPG with $\kappa = |C|$ and lexicographic tie-breaking, the candidates who are lower than the winner in the tie-breaking order are approved by all voters, while those who are higher than the winner in the tie-breaking order are approved by exactly $|V| - 1$ voters.

A variant of Observations 3 and 4 holds for randomised tie-breaking, in cases where equilibrium profiles result in singleton winner sets.

OBSERVATION 5. In a DPG with $\kappa = |C|$ and randomised tie-breaking, if the winning set for an equilibrium profile has size one, then the unique winner is approved by all voters and every other candidate is approved by $|V| - 1$ voters.

By combining these observations, we can see that if the social bonus is uncapped, the structure of the Nash equilibria is very different from what is observed in real-life Doodle polls, as described by Zou et al. [18]. In particular, in every equilibrium where the winner set is a singleton, the winner is approved by all voters and all remaining candidates are approved by (almost) all voters.

Thus, if we define our goal as to model voters' behaviour in Doodle polls, we should focus on the setting where the social bonus is capped, i.e., $\kappa < |C|$; indeed, we consider this case in detail in Section 5. However, the mathematical model of approval voting with uncapped social bonus is of interest per se, as it may be applicable to other approval voting scenarios. Further, the insights obtained by analysing the simpler case of uncapped social bonus will turn out to be helpful for understanding the more complex scenario where the social bonus is capped. Motivated by these considerations, in the rest of this section we analyse the case $\kappa = |C|$ in more detail. While most of our results in this section concern the algorithmic complexity of computing NE, they also provide useful insights into the structure of such equilibria: for instance, we leverage Algorithm 1 to show that, for lexicographic tie-breaking, most profiles do not admit NE.

4.1 Lexicographic Tie-Breaking

When ties are broken lexicographically, it is possible to decide in polynomial time whether there exists a Nash equilibrium (NE) profile where a given alternative wins the election. We denote the corresponding decision problem by $\exists\text{NEWIN}$:

$\exists\text{NEWIN}$: Given a DPG with lexicographic tie-breaking and $\kappa = |C|$, and an alternative $w \in C$, is there a NE with winner w ?

This problem is solved by Algorithm 1, which relies on subroutines provided by Algorithms 2, 3 and 4. The proof that our algorithm is correct relies on the following lemma.

Algorithm 1 $\exists\text{NEWIN}$

Input: DPG $\Gamma = \langle V, 2^C, (U_v)_{v \in V} \rangle$ with $\kappa = |C|$ and tie-breaking order \succ ; candidate $w \in C$.

Output: YES if there exists an equilibrium with winner w ; NO otherwise.

$\text{answer}C^{\text{TOP}} = \text{Test}C^{\text{TOP}}(\Gamma, w)$

if $\text{answer}C^{\text{TOP}} \neq \text{POSSIBLE}$ **then**
 return $\text{answer}C^{\text{TOP}}$

end if

$V' := \{v \in V \mid \delta_v(w) < 1\}$

$\text{answerNonSupporters} = \text{TestNonSupporters}(\Gamma, w, V')$

if $\text{answerNonSupporters} \neq \text{POSSIBLE}$ **then**
 return $\text{answerNonSupporters}$

end if

return $\text{TestDisapprovedAllocation}(\Gamma, w, V')$

LEMMA 1. Consider a DPG with lexicographic tie-breaking and $\kappa = |C|$. Let \mathbf{b} be an equilibrium with winner w . Then for every voter v with $b^v = C$ we have $\delta_v(w) = 1$, while for every voter v and every candidate c with $c \geq_v w$ we have $c \in b^v$.

PROOF. Let v be a voter that approves all candidates in C and assume that $\delta_v(w) < 1$. Then, there exists another candidate c with $\delta_v(c) > \delta_v(w)$. By Observation 3, winner w is approved by all $|V|$ voters. Let C^+ be the set of all candidates that precede w in the tie-breaking order, and let C^- be the set of all candidates that appear after w in the tie-breaking order. By Observation 4, all candidates in C^+ get $|V| - 1$ votes and all candidates in C^- get $|V|$ votes. Suppose that v changes her vote from C to $\{c\}$. If $c \in C^-$, she becomes the only candidate in C with $|V|$ points, so she becomes the unique winner. If $c \in C^+$, then after the change all candidates have at most $|V| - 1$ points and all candidates in $C^+ \setminus \{c\}$ have $|V| - 2$ points, so c wins by the tie-breaking rule. Thus, in both cases c becomes the new winner, so voter v will get a higher utility from the outcome. Moreover, by our assumption about the value of β , his total utility will also increase.

Suppose now that there exist a voter v and a candidate c such that $c \geq_v w$, but $c \notin b^v$. Clearly, approving c is a profitable move for v : the outcome either remains the same or changes from w to c , so his utility from the outcome does not go down, and his social bonus increases by β . \square

Note that it follows from Lemma 1 that no voter may derive zero utility from an equilibrium winner. This, in particular, implies that in games with dichotomous preferences, a candidate can be an equilibrium winner only if she belongs to the top preference level of each voter.

We are now ready to present a sketch of our algorithm; for implementation details, see the pseudocode (Algorithm 1 and Algorithms 2–4). Assume for convenience that the tie-breaking order \succ is given by $c_1 \succ \dots \succ c_{|C|}$.

- First, the algorithm considers the subset of candidates C^{TOP} that consists of the candidates who are at the top preference level of every voter (subroutine $\text{Test}C^{\text{TOP}}$). By Observation 2 and Lemma 1, if $C^{\text{TOP}} \neq \emptyset$ then the election winner belongs to C^{TOP} . Hence, if $C^{\text{TOP}} \neq \emptyset$, our algorithm checks whether w is the tie-breaking winner among the candidates in C^{TOP} . If so, it is possible to construct a NE profile where w wins, by letting all voters vote for their top choices, and also approve some of their less preferred candidates, so that each candidate gets either $|V|$ or $|V| - 1$ votes, depending on her position with respect to w in the tie-breaking order.

Algorithm 2 *TestC^{TOP}* (first stage of \exists NEWIN)

Input: DPG $\Gamma = \langle V, 2^C, (U_v)_{v \in V} \rangle$ with $\kappa = |C|$, tie-breaking order \succ ; candidate $w \in C$.

Output: YES if there exists an equilibrium with winner w ; NO if it is certain that there is no such NE; POSSIBLE otherwise.

```
 $C^{TOP} := \{c \in C \mid \delta_v(c) = 1 \text{ for all } v \in V\}$ 
if  $w \in C^{TOP}$  then
  if  $w \succ w'$  for all  $w' \in C^{TOP}$  then return YES
  else return NO
end if
else
  if  $C^{TOP} \neq \emptyset$  then return NO
  end if
end if
return POSSIBLE
```

Algorithm 3 *TestNonSupporters* (second stage of \exists NEWIN)

Input: DPG $\Gamma = \langle V, 2^C, (U_v)_{v \in V} \rangle$ with $\kappa = |C|$, tie-breaking order \succ ; candidate $w \in C$; the set V' of voters with $\delta_v(w) < 1$.

Output: YES if there exists an equilibrium with winner w ; NO if it is certain that there is no such NE; POSSIBLE otherwise.

```
if  $|V'| > |C|$  then return NO
end if
if  $\exists v \in V'$  such that  $\delta_v(w) = 0$  then return NO
end if
if  $\exists v \in V'$  and  $c \in C$  such that  $\delta_v(c) > \delta_v(w)$  and  $w \succ c$  then
  return NO
end if
return POSSIBLE
```

Hence the algorithm returns YES. Otherwise (i.e., if $w \notin C^{TOP}$, or if w is not the tie-breaking winner among the candidates in C^{TOP}), it returns NO. Finally, if C^{TOP} is empty, the algorithm proceeds to the next step.

- At this stage, the algorithm focuses on the subset of voters $V' = \{v \in V \mid \delta_v(w) < 1\}$ (subroutine *TestNonSupporters*). This set is non-empty, since otherwise our algorithm would have terminated at the previous step. If $|V'| > |C|$ then it is impossible to construct a NE with winner w . This is because, by Lemma 1, in any NE \mathbf{b} with winner w we would have $b^v \neq C$ for each $v \in V'$, so the total number of approvals would be at most $|C| \cdot |V| - |V'| < |C|(|V| - 1)$, but by Observation 4 in NE the score of each alternative is at least $|V| - 1$. Hence, in this case the algorithm returns NO. Otherwise, it moves to the next step.

- Now we can assume that $1 \leq |V'| \leq |C|$. If one of the voters in V' has w among his least preferred alternatives, then by Lemma 1 there will be no NE with winner w , so the algorithm returns NO.

- The algorithm also returns NO if there exists a candidate c with $w \succ c$ and a voter v with $c \geq_v w$. Indeed, by Observation 4, in a NE with winner w , both w and c have $|V|$ votes. Thus, v can make c the winner by changing her vote to $\{c\}$.

- Next, the algorithm decides for each voter $v \in V'$ whether it can allocate him a candidate that he may disapprove (subroutine *TestDisapprovedAllocation*). By Lemma 1, at least one such candidate must exist in a NE, and v must prefer w to this candidate. Let c be the first candidate with respect to the tie-breaking order whom v prefers to w (note that $c \succ w$, since otherwise we would have terminated at the previous step). Then in any NE v disapproves

some candidate c' with $c' \succ c$. Indeed, if b^v contains all candidates c' with $c' \succ c$, then v can make c the winner by changing his vote to $\{c\}$; after this change, everyone has at most $|V| - 1$ point, and the tie-breaking rule favours c over all other candidates. Thus, the algorithm associates v with the prefix of the tie-breaking order that ends just before c ; as argued above, this prefix does not contain w .

- The algorithm now acts greedily, as follows. It orders the voters in V' in the ascending order of the length of the associated prefix of \succ , and considers them in this order. For each $i = 1, \dots, |V'|$, if c_i belongs to the prefix of the i -th voter in this order, the algorithm assigns c_i to that voter; otherwise, it returns NO (note that this happens only if there are more than i voters in V' whose associated prefixes are contained in $c_1 \succ \dots \succ c_i$). When all voters in V' have been processed, the algorithm proceeds to its final step.

- Let C' be the set of candidates that precede w in the tie-breaking order and have not been assigned to voters in V' in the previous step. If $C' = \emptyset$, the algorithm returns YES. Otherwise, for each candidate $c \in C'$ the algorithm seeks a voter that will disapprove c in NE (we need one by Observation 4). To this end, the algorithm checks whether there exists a voter in V that prefers w over c . If such a voter is found for each candidate in C' (we can select the same voter for several candidates in C'), the algorithm returns YES. Otherwise, it returns NO.

Algorithm 4 *TestDisapprovedAllocation* (third (main) stage of \exists NEWIN)

Input: DPG $\Gamma = \langle V, 2^C, (U_v)_{v \in V} \rangle$ with $\kappa = |C|$, tie-breaking order \succ ; candidate $w \in C$; the set V' of voters with $\delta_v(w) < 1$.

Output: YES if there exists an equilibrium with winner w ; NO if it is certain that there is no such NE.

```
for all  $v \in V'$  do
   $k_v := \min\{k \mid \delta_v(c_k) > \delta_v(w)\}$ 
   $D_v := \{c_1, \dots, c_{k_v-1}\}$ 
end for
 $C^{DIS} := \emptyset$ 
repeat
  pick  $v'$  from  $\arg \min_{v \in V'} |D_v|$ 
  if  $D_{v'} = \emptyset$  then return NO
  else
     $c' := \arg \min\{j \mid c_j \in D_{v'}\}$ 
     $C^{DIS} := C^{DIS} \cup \{c'\}$ 
    for all  $v \in V'$  do  $D_v := D_v \setminus \{c'\}$ 
    end for
  end if
   $V' := V' \setminus \{v'\}$ 
until  $V' = \emptyset$ 
 $C' := \{c \in C \mid c \succ w\} \setminus C^{DIS}$ 
for all  $c \in C'$  do
  if  $\exists v \in V$  s.t.  $\delta_v(c) < \delta_v(w)$  then  $C' := C' \setminus \{c\}$ 
  end if
end for
if  $C' := \emptyset$  then return YES
else return NO
end if
```

The main theorem of this section follows easily from the description of the algorithm.

THEOREM 1. *Algorithm 1 solves \exists NEWIN in time polynomial in $|V|$ and $|C|$.*

An easy corollary of Theorem 1 is that we can efficiently check whether a given Doodle poll game has a Nash equilibrium, by

querying Algorithm 1 for each $w \in C$. Further, Algorithm 1 can be used to construct an example of a profile with no Nash equilibrium.

EXAMPLE 1. *Consider an election over a candidate set C with $|C| \geq 3$, where for each candidate $c \in C$ there are two voters v, v' with $\delta_v(c) = \delta_{v'}(c) = 1$, $\delta_v(c') = \delta_{v'}(c') = 0$ for all $c' \in C \setminus \{c\}$; note that all voters are dichotomous. It can be verified that Algorithm 1 will output NO at the second step, which means that the respective Doodle poll game has no pure Nash equilibrium.*

Note that if we add voters to the election constructed in Example 1, the resulting election still has no Nash equilibrium. This illustrates that, under lexicographic tie-breaking, in elections where the number of voters is much larger than the number of candidates, the existence of NE is highly unlikely: essentially, there has to be an alternative that is ranked at the top preference level by (almost) all voters. On the positive side, if there is an alternative that is perfect for all voters, some such alternative wins in every equilibrium.

4.2 Randomised Tie-Breaking

Under randomised tie-breaking, the computational complexity of equilibrium-related problems depends on the number of levels in voters' preferences: we obtain easiness results for dichotomous preferences and hardness results for the general case. We omit the proofs due to space constraints.

We consider the following decision problems.

\exists NE: Does a given DPG with randomised tie-breaking and $\kappa = |C|$ possess a NE?

\exists NESINGLE: Given a DPG with randomised tie-breaking and $\kappa = |C|$, is there a NE where the winning set is a singleton?

\exists NETIE: Given a DPG with randomised tie-breaking and $\kappa = |C|$, is there a NE where the winning set is not a singleton?

For games that have two preference levels, we separately check whether there exist NE with a unique winner, and whether there exist NE with multiple winners. Both problems turn out to be computationally easy, and hence so is \exists NE.

THEOREM 2. *In games with dichotomous preferences, the problems \exists NESINGLE and \exists NETIE (and hence \exists NE) are polynomial-time solvable.*

In contrast, for games with trichotomous preferences, we obtain NP-hardness results for \exists NE and \exists NETIE. Our proof uses the following lemma, which establishes an interesting property of voters' best responses and may therefore be of independent interest.

LEMMA 2. *In a DPG with $\kappa = |C|$, randomised tie-breaking and trichotomous preferences, if in a given profile a voter benefits from approving a candidate at the intermediate preference level, then he benefits from approving all the candidates at that level.*

Using Lemma 2, we obtain the following hardness results.

THEOREM 3. *\exists NE and \exists NETIE are NP-complete for trichotomous preferences.*

The proof of Theorem 3 extends easily to voters with more than three preference levels.

5. CAPPED SOCIAL BONUS

In this section, we consider the variant of our model where κ can be significantly smaller than $|C|$ (in fact, it is enough to assume that $\kappa \leq |C| - 2$). We demonstrate that, in contrast with the case of $\kappa = |C|$ analysed in the previous sections, in this model there are many Nash equilibrium profiles, and we use a variant of trembling hand perfect equilibrium to rule out 'bad' equilibria. It turns out that the voters' behaviour in trembling hand perfect equilibria of our games provides a good match to the behaviour observed in practice in Doodle polls, as described by [18]. Therefore, by capping the social bonus, we can both capture more realistic scenarios and obtain more stable outcomes.

Consider first the Nash equilibria of a DPG with $\kappa \leq |C| - 2$. For any candidate $c \in C$, there exists a Nash equilibrium where c gets $|V|$ approvals, and every other candidate gets κ approvals. Importantly, this holds irrespective of the voters' preferences: c can be universally disliked, and some other candidate c' may be at the top preference level of all voters. This example indicates the need for an equilibrium refinement. To this end, we will now define a modified version of Selten's trembling hand perfect equilibrium [16] in order to apply it in our setting: namely, we assume that the voters' hands tremble independently over each cell in a poll (i.e., a candidate) rather than a whole row (strategy).

Specifically, let $\varepsilon > 0$ be the probability of voter v deviating from his intended action regarding candidate c . We call this quantity the *trembling hand (TH) probability*. Then, the probability that voter v submits a ballot \tilde{b}^v instead of the intended ballot b^v , termed the *ballot TH probability*, is given by

$$P_\varepsilon(\tilde{b}^v | b^v) = \varepsilon^{d(\tilde{b}^v, b^v)} (1 - \varepsilon)^{|C| - d(\tilde{b}^v, b^v)}, \quad (5)$$

where $d(\tilde{b}^v, b^v)$ is the Hamming distance between binary vectors \tilde{b}^v and b^v .

For a subset $S \subseteq V$ of voters and a fixed TH probability ε , the *joint ballot TH probability* is given by the product of individual ballot trembling hand probabilities across the set S :

$$P_\varepsilon(\tilde{\mathbf{b}}^S | \mathbf{b}^S) = \prod_{v \in S} P_\varepsilon(\tilde{b}^v | b^v). \quad (6)$$

For a given voter v , let $-v = V \setminus \{v\}$ denote the set of his opponents in the game. When voter v computes his utility from submitting a ballot b^v against the intended joint ballot \mathbf{b}^{-v} of the other voters, he assumes that other voters' hands (but not his) may tremble independently. Thus, his expected utility for a given TH probability ε is given by the expectation, under joint conditional probabilities as defined by (6):

$$\tilde{U}_\varepsilon(b^v, \mathbf{b}^{-v}) = \sum_{\tilde{\mathbf{b}}^{-v} \in \{0,1\}^{|V|-1}} U(b^v, \tilde{\mathbf{b}}^{-v}) P_\varepsilon(\tilde{\mathbf{b}}^{-v} | \mathbf{b}^{-v}). \quad (7)$$

We call this utility the *expected ε -TH utility* of voter v in (b^v, \mathbf{b}^{-v}) .

Given an $\varepsilon \in (0, 1)$, we say that a ballot b^v is an *ε -TH best response* of voter v to an intended joint ballot \mathbf{b}^{-v} of other voters if it maximises v 's expected ε -TH utility in (b, \mathbf{b}^{-v}) over all possible choices of b . Further, we say that b^v is a *TH best response* of v to \mathbf{b}^{-v} if there exists a threshold ε' such that b^v is an ε -TH best response of v to \mathbf{b}^{-v} for all $\varepsilon \in (0, \varepsilon')$. A ballot profile \mathbf{b} is a *trembling hand perfect equilibrium* if each voter's ballot is a TH best response to other voters' ballots.

To maintain tractability, we restrict our analysis of TH best responses in DPGs with $\kappa \leq |C| - 2$ to the case with lexicographic tie-breaking and dichotomous preferences. While it is plausible that similar results hold for randomised tie-breaking and three or

more preference levels, the results of Section 4 indicate that the analysis for these settings will be significantly more complicated.

Given a voter v with dichotomous preferences, we say that c is *good* for v if $\delta_v(c) = 1$ and *bad* for v otherwise. We show that a voter's TH best response is to approve all of his good alternatives, and also some of the least popular bad alternatives, so that the total number of approved candidates does not exceed κ . Algorithm 5 below details the procedure of constructing a TH best response.

Algorithm 5 TH best response

Input: DPG $\Gamma = \langle V, 2^C, (U_v)_{v \in V} \rangle$ with $\kappa \leq |C| - 2$ and tie-breaking order \succ ; voter $v \in V$; joint strategy $\mathbf{b}^{-v} \in (\{0, 1\}^{|C|})^{|V|-1}$ of voters other than v .

Output: strategy $b^v \in \{0, 1\}^{|C|}$ for voter v .

```

 $C_v^{TOP} := \{c \in C \mid \delta_v(c) = 1\}$ 
 $C_v^{APP} := C_v^{TOP}$ 
 $C_v^{DIS} := C \setminus C_v^{TOP}$ 
for all  $c \in C_v^{APP}$  do  $s_c = s_c(\mathbf{b}^{-v}) + 1$ 
end for
for all  $c \in C_v^{DIS}$  do  $s_c = s_c(\mathbf{b}^{-v})$ 
end for
 $W := \arg \max_{c \in C} s_c$ 
Pick  $w \in W$  so that  $w \succ c$  for all  $c \in W \setminus \{w\}$ 
if  $\delta_v(w) = 1$  then
   $C_v^< := \{c \in C_v^{DIS} \mid s_c + 1 < s_w\}$ 
   $C_v^= := \{c \in C_v^{DIS} \mid s_c + 1 = s_w \wedge w \succ c\}$ 
   $C_v^{SAFE} := C_v^< \cup C_v^=$ 
else
   $C_v^{SAFE} := C_v^{DIS}$ 
end if
repeat
   $C' := \arg \min_{c \in C_v^{SAFE}} s_c$ 
  Pick  $c' \in C'$  so that  $c' \prec c$  for all  $c \in C' \setminus \{c'\}$ 
   $C_v^{APP} := C_v^{APP} \cup \{c'\}$ 
   $C_v^{DIS} := C_v^{DIS} \setminus \{c'\}$ 
   $C_v^{SAFE} := C_v^{SAFE} \setminus \{c'\}$ 
until  $C_v^{SAFE} = \emptyset$  or  $|C_v^{APP}| = \kappa$ 
return  $b^v$ , where  $b_c^v = 1$  for  $c \in C_v^{APP}$ ,  $b_c^v = 0$  for  $c \in C_v^{DIS}$ 

```

First, the algorithm initialises the set of candidates for approval, C_v^{APP} , to be the set C_v^{TOP} of v 's good candidates; the bad candidates are placed in C_v^{DIS} . The algorithm then selects the candidates in C_v^{DIS} that are 'safe' for voter v to approve in addition to his most preferred candidates, under the joint strategy \mathbf{b}^{-v} of the other voters: if the current winner, w , belongs to C_v^{TOP} , then C_v^{SAFE} consists of candidates that would not win the election should v decide to approve them; otherwise, approving any candidate from C_v^{DIS} would not lower the voter v 's utility from the outcome, and so they all are included in C_v^{SAFE} . If C_v^{SAFE} is non-empty and the number of approved candidates does not exceed κ , the algorithm picks the most unpopular alternative from C_v^{SAFE} and moves it from the set of disapproved alternatives, C_v^{DIS} , to the set of approved alternatives, C_v^{APP} . This alternative also gets excluded from C_v^{SAFE} , and the algorithm terminates after C_v^{SAFE} is exhausted or the number of approvals reaches κ .

The following theorem is the main result of this section. It shows that the behaviour of a voter who plays a TH best response is similar to what is observed in practice: a voter approves all of his good candidates and a 'safe' subset of his bad candidates.

THEOREM 4. *Algorithm 5 computes a TH best response for a*

voter $v \in V$ in a given DPG with $\kappa \leq |C| - 2$, lexicographic tie-breaking and dichotomous preferences.

PROOF. Take a voter $v \in V$. By Observation 1, v has a best response where he approves all his good alternatives. However, there may also be other best response strategies where voter v only approves a (sufficiently large) subset of his good alternatives. Now, by the same argument as before, there is also a TH best response where v approves all his good alternatives. Indeed, if a voter disapproves any of his good alternatives, he cannot lower his total utility by adding such an alternative to his ballot, independently of the choices (or mistakes made due to the trembling hand) of his opponents. This is because by doing so he can never lower his utility from the outcome, while his social bonus may only grow. Moreover, in the presence of the trembling hand, it is no longer a best response for v to approve only a subset of C_v^{TOP} , as there is a positive probability that v is a pivotal voter for one of his disapproved good candidates, who will be beaten by v 's bad candidate, thus reducing v 's expected utility. Hence, in a TH best response, a voter must approve all of his good alternatives.

If $|C_v^{TOP}| \geq \kappa$, the algorithm returns C_v^{TOP} as the set of approvals, and this is the only TH best response strategy for voter v . If $|C_v^{TOP}| < \kappa$, it may be beneficial for v to also approve some of his bad alternatives. However, while in the absence of the trembling hand it would be a best response strategy for v to approve any subset of C_v^{SAFE} so that the total number of approved candidates reaches or even exceeds κ , in a TH best response, it becomes unsafe to exceed the cap, as with a positive probability voter v is pivotal for some of his approved bad candidates, while the current winner may be among v 's good candidates. Hence, it is clearly a dominated strategy to approve extraneous bad candidates, i.e., in a TH best response, a voter approves $\min\{|C_v^{SAFE}|, \kappa - |C_v^{TOP}|\}$ of his bad candidates.

It remains to prove that in a TH best response, a voter only approves least popular candidates from C_v^{SAFE} . Assume there are two candidates $c_i, c_j \in C_v^{SAFE}$ such that $s_{c_i}(\mathbf{b}^{-v}) > s_{c_j}(\mathbf{b}^{-v})$. It suffices to show that in the presence of the trembling hand, approving c_i is strictly less beneficial for v than approving c_j , for any combination of approvals over other candidates by voter v . For convenience, we assume that both c_i and c_j are lower in the tie-breaking order than all of v 's good alternatives. The proof for other cases is analogous, but requires a lengthy case analysis.

Note that since $\delta_v(c_i) = \delta_v(c_j) = 0$, if the current winner of the election (that is, the winner under the trembling-hand profile $\tilde{\mathbf{b}}^{-v}$ and the ballot obtained from b^v by removing approvals for c_i and c_j) is not in C_v^{TOP} , then v is indifferent between voting for c_i and voting for c_j . Therefore, we only need to consider the cases where the current winner belongs to C_v^{TOP} . In such cases, the utility from the outcome for voter v can only change (namely, decrease) if his newly approved bad candidate (c_i or c_j) beats the current winner. Thus, we need to show that the probability of c_i beating the current winner is strictly greater than the probability of c_j beating it.

Let z be the number of approvals received by the election winner, and let

$$p_1 = P\left(s_{c_i}(\tilde{\mathbf{b}}^{-v}) = z \mid b^v\right), \quad p_2 = P\left(s_{c_j}(\tilde{\mathbf{b}}^{-v}) \leq z \mid b^v\right).$$

Then, $p_1 p_2$ is the probability that v makes c_i (but not c_j) the election winner. Similarly, let

$$p_3 = P\left(s_{c_j}(\tilde{\mathbf{b}}^{-v}) = z \mid b^v\right), \quad p_4 = P\left(s_{c_i}(\tilde{\mathbf{b}}^{-v}) \leq z \mid b^v\right);$$

then $p_3 p_4$ is the probability that v makes c_j (but not c_i) the election winner. It suffices to prove that $p_1 > p_3$ and $p_2 > p_4$.

To compute p_1 and p_3 , for each $|V|$ -dimensional binary vector with exactly z ones (that is, for each possible $\tilde{\mathbf{b}}^{-v}$ where c_i or c_j gets exactly z votes), we consider its Hamming distance to each of the vectors $\mathbf{b}_{c_i}^{-v}$ and $\mathbf{b}_{c_j}^{-v}$, representing the other voters' intended votes for c_i and c_j , respectively.

For convenience, and without loss of generality, we permute the voters so that $\mathbf{b}_{c_i}^{-v}$ and $\mathbf{b}_{c_j}^{-v}$ are decreasing, that is, we assume that

$$\mathbf{b}_{c_i}^{-v} = (\underbrace{1, \dots, 1}_i, 0, \dots, 0), \quad \mathbf{b}_{c_j}^{-v} = (\underbrace{1, \dots, 1}_j, 0, \dots, 0),$$

where $i = s_{c_i}(\mathbf{b}^{-v})$, $j = s_{c_j}(\mathbf{b}^{-v})$ and $i > j$.

Note that these two vectors are identical in the first j bits. Thus, we can divide all the vectors with z ones into disjoint subsets defined by their configuration in the first j bits. For every vector in some such subset, the first j bits will contribute equally to its distance from $\mathbf{b}_{c_i}^{-v}$ and from $\mathbf{b}_{c_j}^{-v}$, and therefore it suffices to consider the remaining $n = |V| - j$ bits. We will now prove our claim separately for every subset of vectors in our partition.

Fix some such subset, and denote the number of ones in the last n bits of each of its vectors by $y \leq z$ (different configurations of these y ones define different vectors in the subset). Note that $\mathbf{b}_{c_j}^{-v}$ only has zeroes in the last n bits, and $\mathbf{b}_{c_i}^{-v}$ has $t = i - j$ ones and $n - t$ zeroes. The required probabilities would then depend on how many of these y ones fall in the first t bits, and how many of them fall in the remaining $n - t$ bits. We get

$$p_1 = \sum_{k=0}^t \binom{t}{k} \binom{n-t}{y-k} \varepsilon^{y+t-2k} (1-\varepsilon)^{n-y-t+2k} \quad (8)$$

$$p_3 = \sum_{k=0}^t \binom{t}{k} \binom{n-t}{y-k} \varepsilon^y (1-\varepsilon)^{n-y}. \quad (9)$$

In Equation (8), the voters make two types of mistakes: $t - k$ of them consist of placing zeroes instead of ones, and $y - k$ of them consist of placing ones instead of zeroes. Hence, we get $y + t - 2k$ deviations in total. In Equation (9), all y ones are a result of the trembling hand.

For comparison, divide both equations by $\varepsilon^y (1-\varepsilon)^{n-y}$. We have to show that

$$\sum_{k=0}^t \binom{t}{k} \binom{n-t}{y-k} \varepsilon^{t-2k} (1-\varepsilon)^{2k-t} > \sum_{k=0}^t \binom{t}{k} \binom{n-t}{y-k}.$$

Note that $\varepsilon^{t-2k} (1-\varepsilon)^{2k-t} = \left(\frac{1-\varepsilon}{\varepsilon}\right)^{2k-t} > 0$ for any k , so we can omit the terms with $k < \frac{t}{2}$ from the left-hand side. We obtain

$$\sum_{k=\frac{t}{2}}^t \binom{t}{k} \binom{n-t}{y-k} \left(\frac{1-\varepsilon}{\varepsilon}\right)^{2k-t} > \sum_{k=0}^t \binom{t}{k} \binom{n-t}{y-k}.$$

For $k = \frac{t}{2}$ we have $\left(\frac{1-\varepsilon}{\varepsilon}\right)^{2k-t} = \left(\frac{1-\varepsilon}{\varepsilon}\right)^0 = 1$, so we can extract the respective term on both sides of the inequality. By the symmetry of binomial coefficients on the right-hand side, we obtain

$$\sum_{k=\frac{t}{2}+1}^t \binom{t}{k} \binom{n-t}{y-k} \left(\frac{1-\varepsilon}{\varepsilon}\right)^{2k-t} > 2 \sum_{k=\frac{t}{2}+1}^t \binom{t}{k} \binom{n-t}{y-k}.$$

The inequality holds since $\frac{1-\varepsilon}{\varepsilon} > 2$ as $\varepsilon \rightarrow 0$, and $2k - t > 1$.

The proof that $p_2 > p_4$ follows the same lines: we show the inequality for each $0 \leq y \leq z$, and then sum up. \square

Importantly, one can use Algorithm 5 to compute a trembling hand perfect equilibrium of a dichotomous Doodle poll game. To

this end, we start with the profile where each voter approves his good alternatives, and then allow each voter to make a best response move in some fixed order. We obtain the following corollary.

COROLLARY 1. *Given a Doodle poll game with $\kappa \leq |C| - 2$, lexicographic tie-breaking and dichotomous preferences, it is possible to compute a trembling hand perfect equilibrium in time polynomial in $|V|$ and $|C|$.*

6. CONCLUSIONS

Our primary goal in writing this paper was to put forward a plausible explanation to the counterintuitive phenomena observed in Doodle polls. While Zou et al. [18] suggested that these phenomena may be caused by voters' desire to appear cooperative, they stopped short of providing a model of voters' utilities that matches the observed behaviour. Building on their work, we developed a computationally tractable model whose results agree with the real-life data. Besides the basic idea of rewarding voters for approving additional alternatives, our model has two new ingredients: placing a cap on the social bonus and identifying a suitable equilibrium refinement, namely, trembling hand perfect equilibria (THPE).

Our analysis in this paper, together with that of [15] for plurality voting, demonstrates that the concept of trembling hand equilibria is very useful in the context of voting games, where each player (voter) has limited power, which gives rise to multiple (and often undesirable) pure Nash equilibria. It is perhaps not surprising that profiles arising at THPE are intuitively appealing; interestingly, they also turn out to be efficiently computable for a useful class of settings (lexicographic tie-breaking and dichotomous preferences). It may thus be interesting to understand the structure of THPE in other important application scenarios that have received significant attention in algorithmic game theory, such as, e.g., congestion games.

While we see our results for the setting with capped social bonus as our main contribution, the algorithmic results in Sections 4.1 and 4.2 help us identify important features of our model. In particular, we use the algorithm for lexicographic preferences described in Section 4.1 in order to build a profile with no NE and to argue that such profiles are quite likely in large elections. On the other hand, hardness results of Section 4.2 indicate that randomised tie-breaking gives rise to games where finding equilibria is intractable, so, when studying the setting with capped social bonus, we focus on the easier case of lexicographic tie-breaking.

The main take-home message of Zou et al. [18], as well as of our work is that there are settings where voters engage in strategic behaviour under approval voting. In particular, this happens when voters' utilities are different from those assumed in the classic approval model, and specifically take into account social effects. It would be interesting to investigate the potential impact of such effects in other voting scenarios.

A topic that deserves further study is the effects of sequential nature of open polls on voters' behaviour. On the one hand, regardless of the order in which the voters submit their responses, each of them can return to the poll any time and change his strategy. Hence, a single-shot game is indeed a valid model for this scenario. On the other hand, in practice many voters seem to ignore this opportunity and tend to approve more slots the later they join the poll [18]. Hence, sequential effects need closer examination.

Acknowledgments

This work has been supported by the Israeli Center for Research Excellence, Hebrew University, COST Action IC1205 on Computational Social Choice, and ERC StG ACCORD (GA 639945).

REFERENCES

- [1] D. Acemoglu, G. Egorov, and K. Sonin. Equilibrium refinement in dynamic voting games. Technical report, Cambridge, MA : Massachusetts Institute of Technology, Dept. of Economics, 2009.
- [2] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [3] S. J. Brams and P. C. Fishburn. Approval voting. *The American Political Science Review*, 72(3):831–847, 1978.
- [4] S. J. Brams and P. C. Fishburn. *Approval voting*. Springer, New York, 2nd edition, 2007.
- [5] Y. Desmedt and E. Elkind. Equilibria of plurality voting with abstentions. In *ACM EC'12*, pages 347–356, 2010.
- [6] E. Elkind, S. Obraztsova, E. Markakis, and P. Skowron. Equilibria of plurality voting: Lazy and truth-biased voters. In *SAGT'15*, pages 110–122, 2015.
- [7] U. Endriss. Sincerity and manipulation under approval voting. *Theory and Decision*, 74:335–355, 2013.
- [8] K. Etesami, K. A. Hansen, P. B. Miltersen, and T. B. Sørensen. The complexity of approximating a trembling hand perfect equilibrium of a multi-player game in strategic form. In *SAGT'14*, pages 231–243, 2014.
- [9] A. Filos-Ratsikas and P. B. Miltersen. Truthful approximations to range voting. In *WINE'14*, pages 175–188, 2014.
- [10] K. A. Hansen, P. B. Miltersen, and T. B. Sørensen. The computational complexity of trembling hand perfection and other equilibrium refinements. In *SAGT'10*, pages 198–209, 2010.
- [11] M. O. Jackson and X. Tan. Deliberation, disclosure of information, and voting. *Journal of Economic Theory*, 148(1):2–30, 2013.
- [12] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, New York, 1st edition, 1995.
- [13] M. Messner and M. K. Polborn. Miscounts, Duverger's law and Duverger's hypothesis. Technical report, Innocenzo Gasparini Institute for Economic Research, 2011.
- [14] H. Moulin. Dominance solvable voting schemes. *Econometrica*, 47:1337–1351, 1979.
- [15] S. Obraztsova, Z. Rabinovich, E. Elkind, M. Polukarov, and N. Jennings. Trembling hand equilibria of plurality voting. In *IJCAI'16*, pages 440–446, 2016.
- [16] R. Selten. A reexamination of the perfectness concept for equilibrium points in extensive game. *International Journal of Game Theory*, 4:25–55, 1975.
- [17] E. van Damme. *Stability and Perfection of Nash Equilibria*. Springer, New York, 2nd edition, 1991.
- [18] J. Zou, R. Meir, and D. Parkes. Strategic voting behavior in Doodle polls. In *CSCW'15*, pages 464–472, 2015.