How Bad is Selfish Doodle Voting?

Extended Abstract

Barbara M. Anthony
Southwestern University
Georgetown, TX
anthonyb@southwestern.edu

Christine Chung
Connecticut College
New London, CT
cchung@conncoll.edu

ABSTRACT

Doodle polls allow people to schedule meetings or events based on the time preferences of participants. Each participant indicates on a web-based poll form which time slots they find acceptable and a time slot with the most votes is chosen. This is a social choice mechanism known as *approval voting*, in which a standard assumption is that all voters vote *sincerely*—no one votes no on a time slot they prefer to a time slot they have voted yes on. We take a game theoretical approach to understanding what happens in Doodle polls assuming participants vote sincerely. First we characterize Doodle poll instances where sincere pure Nash Equilibria (NE) exist, both under lexicographic tie-breaking and randomized tie-breaking. We then study the quality of such NE voting profiles in Doodle polls, showing that the price of anarchy and price of stability are both unbounded, even when a time slot that many participants vote yes for is selected. Finally, we give some conditions under which the quality of the NE (and strong NE) is good.

KEYWORDS

approval voting; Doodle polls; Nash equilibria; price of anarchy

ACM Reference Format:
Barbara M. Anthony and Christine Chung. 2018. How Bad is Selfish Doodle Voting?. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10–15, 2018,* IFAAMAS, 3 pages.

1 INTRODUCTION AND MODEL

Online scheduling apps such as Doodle (www.doodle.com) are an increasingly popular tool for scheduling meetings and other events. In a Doodle poll, the poll initiator posts a set of possible meeting times, then asks participants to check off the times they are available to meet. This mechanism employed by Doodle for recommending the best time slot is a *social choice function* equivalent to *approval voting*, where each voter in an election must indicate approval or disapproval of each of the candidates. In a Doodle poll, the participants are the "voters" and the time slots are the "candidates." There has been extensive research done in approval voting dating back to the 1970s. For surveys on approval voting from the voting theory literature see [2, 5, 7].

We assume each voter has a privately-held, normalized, utility value (or valuation) for each candidate time slot. To measure the quality of a time slot, we consider the social welfare, or total utility of all voters, for that slot. But while Alrawi et al. [1] study the effect of more "protective" voting behavior compared with more "generous" voting behavior on the social welfare of the winning time slot, this work analyzes the *price of anarchy* and *price of stability* in Doodle polls. The *price of anarchy* (POA) (resp., *price of stability* (POS)) is the worst case ratio, over all possible instances of the game, of the social welfare of an optimal slot to the social welfare of the winning slot(s) at the "worst" (resp., "best") pure NE.

We define a Doodle poll instance to be a triple $I = (A, V, U)$, where $A = \{a_1, a_2, \ldots, a_m\}$ is the set of time slots or *alternatives*, $V = \{v_1, v_2, \ldots, v_n\}$ is the set of voters, and U is the $n \times m$ matrix of utility values $0 \le u_{ij} \le 1$ that each voter $i = 1 \ldots n$ privately holds for each alternative $j = 1 \ldots m$. We say that voter v_i *prefers* alternative a_j to a_k whenever $u_{ij} > u_{ik}$. Given an instance I, we use another $n \times m$ matrix denoted by $R = [r_1, r_2, \ldots, r_n]$ to represent the *voting profile* (or *strategy profile*), where r_i is a binary vector over the m alternatives in A, representing the *vote* or *strategy* of voter i, with $r_i(j) = 1$ (a *yes vote*) if voter v_i approves alternative a_j, and $r_i(j) = 0$ (a *no vote*) otherwise. We consider only *pure strategies* in this work.

Let $s(a_j) = \sum_{i=1}^{n} r_i(a_j)$, or the total count of votes of approval for alternative a_j, be the *score* for an alternative a_j. The default Doodle mechanism (approval voting) chooses the set of one or more *winning alternatives*, W, which maximize the total score, that is $W = \arg\max_{a_j \in A} s(a_j)$. The most commonly-studied tie-breaking rules in the event of multiple alternatives with maximum score ($|W| > 1$) are *lexicographic* tie-breaking, in which the single winning alternative $w \in W$ that comes first in the established tie-breaking order over A is chosen, and *randomized* tie-breaking, which chooses w from W uniformly at random.

A *pure Nash equilibrium* (NE) is a strategy profile where no player can unilaterally *defect* to an alternate strategy (i.e. flip some of their voting bits) such that their payoff strictly increases. We use $OPT(I)$ to denote an *optimal* alternative, one which maximizes the social welfare in a given Doodle poll instance I, and $u(a)$ to denote the total utility (social welfare) of alternative $a \in A$. Hence $OPT(I) = \arg\max_{a_j \in A} \sum_{i=1}^{n} u_{ij}$ and $u(OPT(I)) = \max_{a_j \in A} \sum_{i=1}^{n} u_{ij}$.

As justified in many classical and recent works, e.g.,[2, 3, 6, 8], we assume all voters are *sincere* in their voting, i.e., if $r_i(a_j) = 1$ then $r_i(a_k) = 1$ for all $k \neq j$ where $u_{ik} > u_{ij}$. Let *sincere pure NE* refer to a pure NE where all voters are voting sincerely (and may defect only to sincere strategies) and let $N_s(I)$ denote the set of sincere pure NE for Doodle poll instance I.

Given a Doodle instance I, we define *sincere price of anarchy* $POA(I)$ for that instance to be $\frac{u(OPT(I))}{\min_{R \in N_s(I)} u(w(R))}$, and *sincere price of stability* for an instance I to be $\frac{u(OPT(I))}{\max_{R \in N_s(I)} u(w(R))}$, where $u(w(R))$ is the social welfare of the winning alternative given profile R.

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden.

We define the sincere price of anarchy (resp., stability) of Doodle polls to be the worst-case $POA(I)$: $\max_{I \in \mathcal{I}} POA(I)$, (resp., $POS(I)$: $\max_{I \in \mathcal{I}} POS(I)$), where $\mathcal{I}$ is the set of all Doodle poll instances.

2 SINCERE PURE NASH EQUILIBRIA

Since we analyze price of anarchy and price of stability only over the space of instances that admit sincere pure Nash equilibria, we present our results regarding what these types of instances look like, under both lexicographic and randomized tie-breaking.

Lemma 2.1 (Adapted from [4]). *A voting profile is a sincere pure NE if the two largest scores differ by two or more, under either lexicographic or randomized tie-breaking.*

We refer to an alternative $a_j \in A$ as a *favorite* of voter v_i if $u_{ij} \geq u_{ik}$ for all $k \neq j$. And an alternative $a_j \in A$ is a *kth favorite* of voter v_i if there are exactly $k-1$ alternatives j' for which $u_{ij'} > u_{ij}$.

Corollary 2.2. *If two or more voters have a favorite alternative in common, then there is a sincere pure NE where the set of winning alternatives W consists precisely of that favorite alternative, under both lexicographic and randomized tie-breaking.*

Corollary 2.3. *If the number of voters exceeds the number of alternatives, that is, $n > m$, then there is a sincere pure NE, under both lexicographic and randomized tie-breaking.*

The above results collectively describe a rather large space of instances where a sincere pure NE always exists. Furthermore, the following lemma ensures the existence of a sincere pure NE when $n = m$.

Lemma 2.4. *If the number of voters equals the number of alternatives, that is, $n = m$, then there is a sincere pure NE, under lexicographic tie-breaking.*

These conditions greatly limit the potential instances which may not have a sincere pure NE. However, we have found that sincere pure NE do not always exist, under either lexicographic or randomized tie-breaking.

Theorem 2.5. *Sincere pure NE do not always exist in Doodle polls under lexicographic tie-breaking.*

We now provide a broad categorization of instances which do have sincere pure NE under randomized tie-breaking, but again show that sincere pure NE do not always exist. We say an $n \times n$ instance does not have *distinct ith favorites* if for $i \in 1, 2, \ldots, n-1$, some alternative is the ith favorite of two or more voters.

Lemma 2.6. *If an $n \times n$ instance does not have distinct ith favorites, then it has a sincere pure NE under randomized tie-breaking.*

Lemma 2.6 together with Corollary 2.3 indicates that the space of instances under consideration in this work, those that admit sincere pure NE, is quite general and large, and in the case of $n = m$, only instances that meet the strict structural requirement of distinct ith favorites do not have sincere pure NE.

Theorem 2.7. *Sincere pure NE do not always exist in Doodle polls under randomized tie-breaking.*

Finally, we find that both price of anarchy and price of stability are unbounded, regardless of which tie-breaking mechanism is used.

Theorem 2.8. *The sincere price of anarchy is unbounded in Doodle polls, under both lexicographic and randomized tie-breaking.*

Theorem 2.9. *Sincere price of stability is unbounded in Doodle polls, even when $|OPT| \approx n-2$. Furthermore, the claim holds under both randomized and lexicographic tie-breaking.*

3 BOUNDS ON POS AND STRONG POA

In this section we describe some situations where price of stability is good. Since Corollary 2.2 guarantees that there is a NE which selects an optimal alternative, the following additional corollaries identify situations in which POS is 1.

Corollary 3.1. *In a Doodle poll instance I, if an optimal alternative is a favorite of two or more voters, then $POS(I) = 1$.*

Corollary 3.2. *If there are two or more 'indifferent' voters with identical valuations on all alternatives in a Doodle poll instance I, then $POS(I) = 1$.*

We also provide the following characterization of the set of Doodle polls instances where the expected social welfare in the best NE is optimal.

Theorem 3.3. *Given a Doodle poll instance I under randomized tie-breaking, $POS(I) = 1$ if and only if there is no alternative that $n-1$ voters prefer to an optimal alternative. I.e., for each non-optimal alternative, at least 2 players prefer an optimal alternative to it.*

A *strong NE* is a voting profile where no subset (or "coalition") of voters can all simultaneously defect and improve their payoff. All strong NE are NE, and strong NE may not always exist. The *strong POA* (resp., *POS*) is defined as the ratio of the total utility of an optimal alternative to the total utility of the alternative chosen in the worst (resp., best) strong NE, assuming one exists.

We now state our result that strong POA is at most 4 when there is an alternative with utility at least $\frac{3n}{4}$, or more generally:

Theorem 3.4. *Given a Doodle poll instance $I = (A, V, U)$ that admits a strong NE, if $u(a_j) \geq \rho n$ for some $a_j \in A$, $1 \geq \rho > 1/2$, then strong $POS(I) \leq$ strong $POA(I) \leq 1/(\rho - 1/2)$, which approaches 2 as ρ approaches 1.*

4 CONCLUSION

Our results have shown that there are many natural Doodle instances that admit sincere pure Nash Equilibria. In particular, almost all instances where the number of voters is at least the number of candidates (that is, $n \geq m$) admit sincere pure NE under both randomized and lexicographic tie-breaking. It remains future work to determine when sincere pure NE exist in the case where $m > n$ (the number of candidates exceeds the number of voters), which is not common in standard approval voting settings, but is not so unusual to encounter in a Doodle poll.

While the price of anarchy and price of stability are both unbounded, the conditions we found that give rise to these cases seem rather particular and unlucky. We also show that there is also a large set of realistic Doodle instances where $POS = 1$; for example, $POS = 1$ when the optimal time slot is the favorite of at least two voters. Finally, we also show that strong POA is reasonable when there is at least one time slot with total utility more than $n/2$.

REFERENCES

[1] Danya Alrawi, Barbara M. Anthony, and Christine Chung. 2016. How Well Do Doodle Polls Do?. In *Social Informatics - 8th International Conference, SocInfo 2016, Bellevue, WA, USA, November 11-14, 2016, Proceedings, Part I (Lecture Notes in Computer Science)*, Emma S. Spiro and Yong-Yeol Ahn (Eds.), Vol. 10046. 3–23. https://doi.org/10.1007/978-3-319-47880-7_1

[2] Steven J. Brams and Peter C. Fishburn. 1983. *Approval Voting*. Birkhauser Boston. https://books.google.com/books?id=YIuGAAAAMAAJ

[3] Steven J. Brams and M. Remzi Sanver. 2006. Critical strategies under approval voting: Who gets ruled in and ruled out. *Electoral Studies* 25, 2 (Jun 2006), 287–305. https://doi.org/10.1016/j.electstud.2005.05.007

[4] Jean-François Laslier and M. Remzi Sanver. 2010. The Basic Approval Voting Game. In *Handbook on Approval Voting (Studies in Choice and Welfare)*, Jean-François Laslier and M. Remzi Sanver (Eds.). Springer, Berlin, Heidelberg, 153–163.

[5] Jean-François Laslier and M. Remzi Sanver. 2010. *Handbook on Approval Voting*. Springer, Berlin, Heidelberg. https://books.google.com/books?id=mQBEAAAAQBAJ

[6] Svetlana Obraztsova, Maria Polukarov, Zinovi Rabinovich, and Edith Elkind. 2017. Doodle Poll Games. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, Sao Paulo, Brazil (AAMAS '17)*. 876–884. http://dl.acm.org/citation.cfm?id=3091125.3091249

[7] Robert J. Weber. 1995. Approval voting. *The Journal of Economic Perspectives* 9, 1 (1995), 39–49.

[8] James Zou, Reshef Meir, and David Parkes. 2015. Strategic Voting Behavior in Doodle Polls. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work (CSCW '15)*. ACM, New York, NY, USA, 464–472. https://doi.org/10.1145/2675133.2675273