

Approval Voting Behavior in Doodle Polls

James Zou, Reshef Meir and David Parkes

Abstract

Doodle is a simple and popular online system for scheduling events. It is an implementation of the approval voting mechanism, where candidates are the time slots and each responder approves a subset of the slots. We analyze all the Doodle polls created in the US from July-September 2011 (over 340,000 polls), consisting of both hidden polls (where you cannot see other people's votes) and open polls (where you can see all the previous responses). By analyzing the differences in behavior in hidden and open polls, we gain unique insights into strategies that people apply in natural voting settings. Responders in open polls are more likely to approve slots that are very popular or very unpopular, but not intermediate slots. We show that this behavior is inconsistent with models that have been proposed in the voting literature, and propose a new model based on combining personal and social utilities to explain the data.

1 Introduction

Doodle is a simple and popular online system for scheduling. After the poll initiator defines the possible dates and time slots (say, for an upcoming event), each participant can mark several time slots as available. The poll initiator then uses the reported availability to determine the event time.

One way to think about Doodle polls is as a *voting mechanism*. That is, after the possible slots have been announced, each participant is a voter, with preferences over event times. These preferences may be derived from constraints and considerations such as convenience of commuting, other events on the same day, and so forth. By marking some of the slots, a participant discloses partial information about her preferences, which is then aggregated with the votes of others. If we assume that the option with most votes is selected, then we have a well-defined *social choice function*, mapping any set of votes to an outcome. When there are no constraints on which or how many slots can be marked by each participant, which is typically the case on Doodle, this social choice function is known as the *Approval voting system*.

Approval is a special type of voting rule. On one hand, it is very simple in the sense that voters do not need to specify a full preference order. On the other hand, even if we accept the main axiom of social choice that a transitive (strict or weak) preference order exists, it is not clear how a participant should vote even if she wants to reveal her truthful preferences. Which subset of candidate times should be approved? On top of this, a participant may take into account various strategic and social considerations when casting her vote. For example, she may hide some of her available slots in hope that another, more convenient slot will be selected. Alternatively, she may mark a less convenient slot if she believes that this would enable more participants to attend.

Doodle polls provide a unique opportunity to study the range of strategies that people apply in such voting situations, due to three main reasons. First, there is a huge database of polls that can be used for analysis. In this paper we used more than 14 million votes from 2 million responders in over 340,000 polls. For comparison, typical datasets used in the social choice literature contain few independent polls or elections (sometimes just one), where the number of responders in each poll ranges from dozens to thousands.¹ See, e.g., [7, 17, 21].

Second, the database consists of both *hidden polls*, where participants do not see the other votes, and *open polls*, where participants can see the votes of previous participants. By using the hidden-poll database as a baseline, we can study how voting behavior is affected by the information available to the participant and avoid other confounds.

¹Larger datasets of *preferences* also exist [15], but these are collected in settings where there are no incentives involved.

Third, Approval voting has been extensively studied in the social choice and AI literature, and there are many models of both truthful and strategic voting. The Doodle data can be used to test the assumptions underlying these models in a particular real scenario, as well as their predictions on voting outcomes.

1.1 Summary of Key Findings

1. The average reported availability is higher in open polls compared to hidden polls.
2. In both open and hidden polls, there is a decline in reported availability over time. The relative rate of decline is similar between open and hidden polls.
3. Responses in open polls have higher positive correlation with previous responses compared to hidden polls.
4. Open polls have higher response rates for very popular and unpopular time slots. Intermediate time slots have similar response rates between open and hidden polls.
5. These empirical results are inconsistent with traditional models of approval voting. We propose a new model, whereby responders in open polls vote for their preferred time slots while also trying to appear cooperative, to explain the data.

1.2 Related work

Voting patterns in Doodle have been studied by Reinecke et al. [19], who focus on cultural differences between countries. In particular they show that participants in collectivist countries tend to coordinate more with one another in open polls. We used data from a single country (US) to avoid such differences as much as possible.

The phenomenon of vote coordination in open online polls is related to *herding* [5, 24, 12, 1], where information revealed in early votes influences voting dynamics and leads to a failure of information aggregation. Herding typically refers to situations where a voter faces a binary choice (e.g., to recommend a product or not, or choosing between two candidates).

The implications of sequential voting in more complex scenarios has also been studied [2, 9]. However these models deal with equilibrium analysis, thereby assuming very high sophistication of the voters. They also do not consider the Approval system.

Brams and Fishburn [3] offered the first systematic model of strategic behavior under Approval voting. They assume that voters have a weak transitive preference order over alternatives, and define a vote as *sincere* if the voter prefers any approved alternative to every other alternative. Crucially, a voter may have more than one sincere way to vote. They prove that a voter who has at most three levels of preference² is always better off by voting sincerely. However, a voter may prefer to vote insincerely if she has four or more levels of preference.

A key issue in strategic voting models is the assumption about what information voters have when deciding on their vote. Brams and Fishburn assume that voters have no knowledge whatsoever about the preferences or actions of others, and restrict themselves to notions based on dominance. A more elaborate model for Approval voting was studied by Weber [23], based on the general voting theory of Myerson and Weber [16]. Here voters have a common prior distribution over the total number of votes obtained by each alternative, and they each try to maximize their expected utility w.r.t. this distribution. Weber shows that the optimal vote (i.e., the rational best response of a voter to any distribution) is always sincere.

While the Weber model assumes that voters are highly sophisticated agents capable of probabilistic calculations, a heuristic strategy called the *leader rule* was suggested by Laslier [11]. This

²That is, alternatives can be partitioned to three sets, and the voter is indifferent between the alternative in each set.

is a simple and sincere strategy (see details in the next sections), and Laslier was able to show that it is optimal for a self-interested, rational voter in a special case of the Weber model in which there is a common prior on the rank of alternatives.

Strategic behavior in Approval voting was also studied within the AI and ComSoc communities, with some researchers focusing on the design of variations that cope better with such behavior [22, 10, 6, 20].

Other researchers have studied social factors affecting voting behavior, albeit not under the Approval system [8, 13]. In particular, the MBD model, named after Manski [14], Brock and Durlauf [4], assumes that a voter tries to meet the expectations of her peers (see [13]). The MBD model completely ignores the self-interest aspect of voting that is standard in the AI literature: the alternative that actually gets selected does not factor into the utility of the voter. The model that we propose and support with the Doodle poll data combines social and self-interest considerations.

2 Preliminaries

Let $A = \{a_1, a_2, \dots, a_M\}$ denote the time slots designated by the initiator of the poll. Slots are also referred to as *candidates* or *alternatives*. We denote the responders (or *voters*) by $V = v_1, v_2, \dots, v_N$, where voters are in temporal order so that v_n is the n 'th responder to the poll. The *response* (or *vote*) of v_n is the set of slots that she approves, denoted $r_n \subseteq A$. We also think of r_n as a binary vector, with $r_n(a) = 1$ if she approves slot a and $r_n(a) = 0$ otherwise.

For $A' \subseteq A$, and voter v_n , let $r_n(A')$ denote the average of $r_n(a)$ for alternatives $a \in A'$. In particular, we denote by $\bar{r}_n = r_n(A)$ the fraction of alternatives that v_n approved. For a set of votes R , we denote by $s(a, R) = \sum_{r \in R} r(a)$ the *score* of alternative a , aggregating all votes in R . We denote by $\mathbf{s}(R) = (s(a, R))_{a \in A}$ the score vector over R . The *winner* of a poll is assumed to be $w(R) = \operatorname{argmax}_{a \in A} s(a, R)$, where we omit R when it is clear from the context. The tie breaking will not matter for our purposes.

Let $R_{\leq n} = (r_i)_{i \leq n}$ denote the collection of all votes by voters up to and including v_n . We adopt $s_{\leq n}(a)$ and $\mathbf{s}_{\leq n}$ as a shorthand for $s(a, R_{\leq n})$ and $\mathbf{s}(R_{\leq n})$, respectively. Given the first n voters, we divide the time slots into three sets $\{\text{Popular}_{\leq n}, \text{Intermediate}_{\leq n}, \text{Unpopular}_{\leq n}\}$, so that the slots with the highest third, middle third, and lowest third of $\mathbf{s}_{\leq n}$ are in $\text{Popular}_{\leq n}$, $\text{Intermediate}_{\leq n}$ and $\text{Unpopular}_{\leq n}$, respectively, breaking ties at random.

2.1 Hypotheses

In open polls, responder v_n observes all the previous responses $R_{\leq n}$, and v_n 's response r_n is public to all other responders. In contrast, in hidden polls v_n does not see any other response and r_n is only known to the poll organizer. Our goal is to investigate how responder's behavior changes when all the responses are public. We state several hypotheses, and later test whether they are corroborated by the data.

Our first hypothesis is bidirectional, as it is not a-priori clear whether seeing previous responses would make the next responder approve more slots.

[H 1.1]. *The fraction of time slots approved by voters differs between hidden and open polls.*

We hypothesize that Doodle responders, on average, want to find time slots that work for other people.

[H 1.2]. *The correlation of the vote r_n with the aggregated previous votes $\mathbf{s}_{\leq n-1}$ is positive and is higher in open polls than in hidden polls.*

We also conjecture that participants in open polls will tend to vote more for popular alternatives, and will be more reluctant to approve unpopular alternatives:

[H 1.3]. The probability that a voter will approve a popular alternative, $r_n(\text{Popular}_{\leq n-1})$, is higher in open polls than in hidden polls.

[H 1.4]. The probability that a voter will approve an unpopular alternative, $r_n(\text{Unpopular}_{\leq n-1})$, is lower in open polls than in hidden polls.

3 Empirical results

3.1 Data collection

We obtained all the polls created by US users on Doodle over the three month period July-September, 2011. We focus our analysis on polls with at least three participants, at least four time slots and only yes/no options (and thus approval polls). There are two types of these yes/no polls: *open* and *hidden*. In open polls, responders can see all the previous responses. In hidden polls, a responder does not see any previous responses. We had 345,297 open polls and 7,390 hidden polls that passed the filtering.

Key assumption. In our analysis, we compare certain statistical properties of the open polls with that of the hidden polls. We do this to isolate the effect of observations of previous responses and to avoid confounds. In order for this comparison to be reasonable, we need open and hidden polls to be generated from similar distributions of activities or events, so that the main difference in response patterns are driven by the fact that responders see previous responses in open polls. Our collection of polls has a median of 5 responders and 12 time slots. All of the patterns we discuss below are robust if we stratify and compare open and hidden polls with the same number of responders and time slots.

3.2 Average availability

We define the *aggregate availability* of a poll to be the average proportion of 'yes' votes from all responders, $\text{avg}(\bar{r}_n)_{n=1}^N$. Open polls have average availability of 0.53, which is significantly higher than the 0.39 average availability of hidden polls ($p < 10^{-5}$). After controlling for the number of time slots and participants in polls as covariates in linear regression, we obtain similarly significant increase in availability in open polls. In addition, we compute the average availability for each time slot, i.e. each column of R , aggregated over all polls. We observe uniform availability across the time slots, showing no bias in approvals for leftmost or rightmost polls. Moreover, for each time slot, open polls have higher availability than hidden polls. We conclude that Hypothesis 1.1 is supported, in the positive direction.

We can also measure the availability for the first responder, second responder, and so forth. The availability of an individual is the fraction of time slots approved, \bar{r}_n . In both open and hidden polls, we observe that the average availability declines monotonically for later responders. This decline could be due to responders of the poll becoming more constrained the more they wait to answer the poll.

Both open and hidden polls have similar rate of decline. We see this in Figure 1. This plots the normalized average availability for responders 2 through 10, in each of open and hidden polls. We ignore responder 1 since this person is likely to be the poll creator and an outlier. Indeed, responder 1 shows 12% and 20% higher availability than the poll average in hidden and open polls, respectively. For the purpose of this plot, we normalize availability by dividing the availability by responder 2's availability.

3.3 Correlation with previous responders

For each responder v_n we compute the Pearson correlation between the response vector r_n and $s_{\leq n-1}$, the sum over all previous responses. This measures how likely v_n is to agree with the

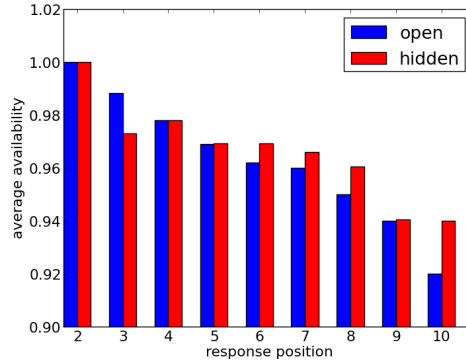


Figure 1: Normalized average availability as a function of responder position in hidden and open polls.

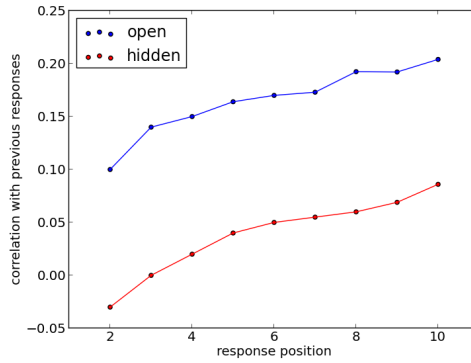


Figure 2: Correlation with previous responses in hidden and open polls.

previous responses.

In both open and hidden polls, the correlation increases monotonically in response position (see Figure 2). Open polls have significantly higher correlation than hidden polls, supporting Hypothesis 1.2. This suggests that later participants actively try to match the previous responses that they observe. As a control, we randomize the open and hidden polls by permuting each row of a poll. The randomized polls, as expected, have zero correlation for all the response positions. The small positive correlation observed in hidden polls could be due to that a subset of slots are intrinsically more popular among all responders, and they are also more likely to be selected by later responders who have additional constraints.

3.4 Where goes the extra availability?

To understand why there is higher availability in open polls, we investigate the distribution of votes for different types of time slots in open and hidden polls. We compute the response curves for responder in position $n = 11$, focusing on polls with ≥ 11 responders (32,527 open polls and 2434 hidden polls). The response curve is computed by looking at the response of the 11th responder, averaged across all polls. See Figure 3. The dots show the probability that the 11th responder inputs

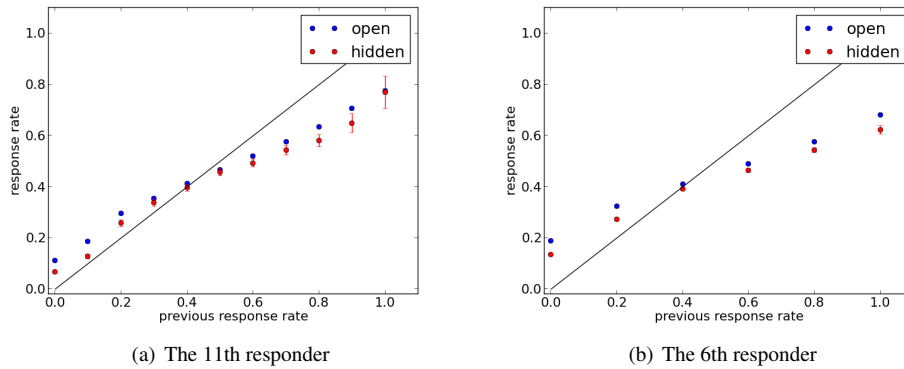


Figure 3: **Response curves for the 11th and 6th responder.** In the left figure, each dot represents the probability that the 11th responder approves a slot (y-axis) if the slot has been approved by i out of the first 10 responders (x-axis), for $i = 0, \dots, 10$. Error bars shows the 95% confidence interval. For open polls, the error bars are very small. The right figure presents the same data for the 6th responder.

‘available’ (y-axis) conditioned on i out of the first 10 votes ‘available’ for $i = 0, 1, \dots, 10$ (x-axis). Both curves are S-shaped due to mean-reversion. Open polls show higher response rate for time slots with both low availability and high availability. Interestingly, open and hidden polls show similar response rates for time slots with intermediate popularity (4/10 and 5/10). To additionally verify this result, we also compute the response curves at $n = 6$ (see Figure 3 on right). Again, open polls exhibit consistently higher response rate compared to hidden polls, with the smallest gap for the intermediate time slot with 2/5 popularity. The difference in response rates for the intermediate time slots between hidden and open polls is not statistically significant.

Thus the Doodle data supports Hypothesis 1.3, and rejects Hypothesis 1.4. In fact, H1.4 is *inverted*, as we observe significantly higher response rates ($p < 10^{-5}$) for the least popular time slots in open polls than in hidden polls.

3.5 Conditional response rate

The response curves in Figure 3 are aggregated over all voters. Two possible explanations to the observed patterns are: (a) responders choose to mark both highly popular and highly unpopular slots; or (b) there are two types of voters, where some tend to mark the popular slots, and others tend to mark the unpopular ones. We compute conditional response rates and show that it support the first explanation.

Conditioned on the $n = 11$ individual having approved a 10/10 time slot (i.e., a very popular one), we compute her response rate for other time slots. In particular, we are interested in her conditional response rate for other popular (9/10, 8/10), intermediate (5/10, 4/10) and unpopular (1/10, 0/10) time slots. See Table 1. Comparing the conditional response rates between open and hidden polls reveals several surprises:

1. In open polls, people who have voted for a 10/10 slot are significantly more likely to approve unpopular slots (0/10 and 1/10) than in hidden polls.
2. In open polls, people who have voted for a 10/10 slot are significantly less likely to approve an intermediate (5/10) slot than in hidden polls.

Table 1: Conditional response rates for individuals who have approved a 10/10 slot.

	0/10	1/10	4/10	5/10	8/10	9/10
open	0.08	0.20	0.44	0.51	0.70	0.76
hidden	0.05	0.09	0.55	0.66	0.69	0.76

Table 2: Conditional response rates for individuals who have approved a 9/10 slot.

	0/10	1/10	4/10	5/10	8/10	10/10
open	0.14	0.24	0.46	0.51	0.69	0.81
hidden	0.03	0.23	0.50	0.56	0.68	0.83

3. In contrast, the conditional response rates for other popular time slots are similar between open and hidden polls.

To further validate these results, we also compute the conditional response rates for the $n = 11$ individual who has approved a 9/10 time slot, and observe similar patterns (Table 2). This means that the same responders who approve popular slots are in effect shifting some of their votes from medium to unpopular slots.

4 Analysis

The pattern observed in Figure 3 and Tables 1, 2 suggests that in open polls responders resort to *strategic* behavior, whereby they do not simply mark all the time slots that they are available for. In order to analyze possible incentives for such behavior, we introduce preference models and concepts commonly used in the theory of Approval voting.

Let P^n denote the preference of voter v_n . P^n is a partition of $\{a_1, \dots, a_M\}$ into sets $\{A_1^n, \dots, A_K^n\}$ such that voter v_n is indifferent about slots in the same set A_k^n , and strictly prefers any slot in A_k^n to any slot in $A_{k'}^n \forall k < k'$. We allow the possibility that a set is empty. We denote $a \succeq_n a'$ if $a \in A_k^n, a' \in A_{k'}^n$, and $k \leq k'$; $a > a'$ if $a \succeq a'$ and not $a' \succeq a$.

The K levels can be thought of as *preference levels*. A responder with $K = 2$ is called *dichotomous*, and is simply “available” at some times and “unavailable” at others. A responder with $K = 3$ may distinguish between times that are convenient, available but inconvenient and not available, and so on. W.l.o.g., since levels can be empty, we assume that all the responders have the same number of levels K .

A vote r_n is *sincere* if for all $a \in r_n, a' \notin r_n$, it holds that $a \succeq a'$. A vote r'_n *dominates* vote r_n , if for any set of votes R that do not include a vote from voter v_n , voter v_n weakly prefers $w(R \cup \{r'_n\})$ to $w(R \cup \{r_n\})$, with a strict preference for at least one set R . A vote r_n is *admissible* if there is no vote r'_n that dominates r_n . Intuitively, if a voter has fixed preference levels and is only interested in the outcome of the vote (that is, in which slot is selected), then she will always submit admissible votes.

Proposition 1 (Brams and Fishburn [3]). *A vote r_n in an Approval voting system is admissible if and only if $A_1^n \subseteq r_n$ and $A_K^n \cap r_n = \emptyset$.*

In words, a vote is admissible if and only if a voter at least votes on all top-level choices and no bottom-level choices. Given this preference level viewpoint, the behavior of responder v_n can be described by a *response function* f . This function takes as input the preferences P_n and the available information on previous responses $s_{\leq n-1}$. The output of response function f is a subset of time slots r_n , taken to be the vote, and the function f may or may not be deterministic. We say that a response function f is sincere if it generates a sincere vote. Different voters may apply different response

functions, but ideally we would like to be able to explain voting behavior using a small number of simple response functions, or even a single response function.

4.1 Testing response functions

We describe some response functions that have been suggested in the literature, and use the Doodle dataset to refute each one in our setting. We assume that in hidden polls, responders submit admissible votes consistent with Prop. 1.

Random cutoff. The simplest response function is to mark the most preferred $q \geq 0$ candidates, where q is fixed or is sampled from some distribution. We call such a response function *q-cutoff*. In the *random cutoff*, q is sampled uniformly from $\{0, 1, 2, \dots, M\}$.

Some variations of the random cutoff function have been studied in [18]. They also describe a more general class of response functions which they call *size independent*. Crucially, all size independent functions only use the preferences P_n as input, and all are sincere. Another variation is to choose an availability level $q > 0$, and approve all slots in A_k^n , $k < q$, plus a random subset in level A_q^n . We refer to this variation as the *q-level cutoff function*. Irrespective of the particular variation, the responder does not use any information based on votes from previous responders. Hence it cannot explain the difference that we observe between open and hidden polls.

Lazy responses. Determining one’s own availability for a certain time slot may require some effort. Recognizing this, we consider a simple variation of the cutoff model, where responders refrain from even checking their own availability for time slots that do not seem like they are going to be selected. Thus a *lazy q-level cutoff* function only considers responding to plausible time slots (each responder might have her own threshold for what level of support makes a slot plausible). A lazy response function requires access to $s_{\leq n-1}$. Lazy responses may not be sincere, and may not even be admissible if the voter refrains from assessing implausible slots in A_1^n . In the lazy model, conditioning on voting for a Popular time slot, we would expect the responder to be less likely to vote for an Unpopular slot in open polls than in hidden polls. Empirically we observe the opposite, and thus reject this model.

The Leader rule. Let x and y be the *leader* and the *challenger* after the first $n - 1$ votes. That is, $x = \operatorname{argmax}_{a \in A} s_{\leq n-1}(a)$, $y = \operatorname{argmax}_{a \neq x} s_{\leq n-1}(a)$. In case of tie, a leader and a challenger is randomly selected from the set of candidates with the most votes. The Leader Rule [11] stipulates that voter v_n will approve all candidates that she strictly prefers to x , and will approve x if and only if x is strictly preferred to y . It is easy to see that the Leader rule is sincere.

To test the Leader Rule, we generate synthetic open polls using a simple preference model constructed from the hidden poll responses. For each hidden poll responder, we assume all of her m approved slots are in level 1 (most preferred); this is consistent with Prop. 1. Then we assign a random preference ordering (starting from $m + 1$ onward) to the other slots. Thus for each responder, we obtain a preference ordering over all time slots. Then we simulate an open poll by defining responses based on the Leader Rule.

Figure 4 (left) shows the $n = 11$ response curve for these synthetic open polls. It exhibits uniformly higher response rates in open polls across all slots. In particular simulated open polls show significantly higher ($p < 0.01$) response rates for Intermediate slots. We conclude that the Leader Rule does not replicate the pattern in Figure 3, where only Popular and Unpopular slots get additional votes and Intermediate slots show statistically similar response rate between open and hidden polls. Therefore we reject this model.

5 A Social Voting Model

Social utility. In many scenarios there are factors that affect the utility of the voter other than just the outcome. For example, voters may want to act in accordance to the expectations of their peers, as in the MBD model [14, 4]. Also in open Doodle polls, responders' names are made public in addition to their votes, so that they may want to approve more time slots to appear cooperative. Social utility is one possible explanation for approving Popular slots. But why would a responder mark an Unpopular slot? We conjecture that there is an implicit social expectation that every responder will mark as many slots as possible. Therefore a responder in open polls may be motivated to mark more slots. While any additional slot marked increases the social utility, marking inconvenient slots is risky, since these slots may come out as winners in the end. The risk is lower if the additional inconvenient slots are very unpopular, as these would be unlikely to win.

Motivated by this consideration, we propose the following **Social Voting** response function for a voter v_n . For simplicity, we consider three level preferences, $K = 3$.

- Approve A_1^n .
- Approve $A_2^n \cap \text{Popular}$.
- Approve $A_2^n \cap \text{Unpopular}$.

According to this model, the slots not approved by the voter are the second level, intermediate popularity slots, and also the least preferred slots.

Proposition 2. *Social voting is admissible and (for $K = 3$) sincere.*

The proof follows immediately from [3]. Our theory is that the Unpopular slots are marked to appear cooperative with the social norm, and since they have a low chance of getting selected anyways. A variation of **Social Voting** is for the responder to also approve $A_3^n \cap \text{Unpopular}$. In our simulation framework (below), this leads to much higher response rates for Unpopular slots than observed in real open polls, hence we reject this alternative.

Simulation Results We show that our proposed behavioral model produces response curves patterns qualitatively similar to that observed in the real data. As in testing the Leader Rule, we take the real hidden polls and generate synthetic open polls based on this behavioral model. For each hidden poll, we assume that a responder v_n approves A_1^n . To generate a responder \hat{v}^n for the corresponding synthetic open poll, we set: $\hat{A}_1^n = A_1^n$; \hat{A}_2^n to be a random subset among the time slots *not* approved by v_n ; and \hat{A}_3^n to be the remaining slots. We choose the size of the random subsets \hat{A}_2^n so that the availability averaged over all the synthetic open polls matches the average availability of the real open polls. Given these assumed preference levels, each voter v_n in a synthetic open poll partitions the time slots into $\text{Popular}_{\leq n-1}$, $\text{Intermediate}_{\leq n-1}$ and $\text{Unpopular}_{\leq n-1}$ and approves $\hat{A}_1^n \cup (\hat{A}_2^n \cap (\text{Popular}_{\leq n-1} \cup \text{Unpopular}_{\leq n-1}))$.

Figure 4 shows the $n = 11$ response curve for these synthetic open polls under this behavior model, combined with the curve for real hidden polls. It matches the pattern observed in the real data (see Figure 3). The synthetic open polls have significantly higher ($p < 10^{-5}$) response rates for popular and unpopular slots, and there is no statistically significant difference with hidden polls for intermediate slots.

6 Discussion

Our analysis shows that responders are more likely to approve highly popular and unpopular time slots in open polls than in hidden polls. When a popular slot emerges, a responder might feel the

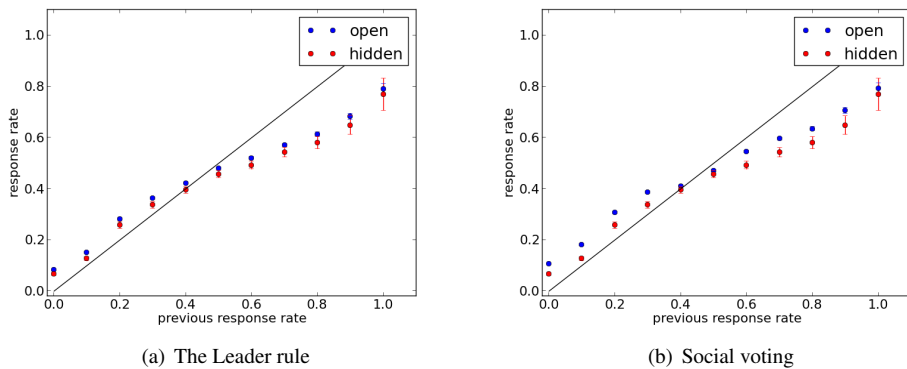


Figure 4: Response curve for synthetic open polls generated by the Leader rule (left) and by the Social voting rule (right). Hidden polls are from the real data.

need to approve it to be cooperative or because she truly wants a time that works for most people. Moreover, because votes are public in open polls, there might be social pressure for responders to mark as many slots as possible to appear flexible. In this case, the ‘safe’ strategy is to vote for unpopular slots that are unlikely to win in addition to her preferred slots. We propose a Social Voting model that captures this phenomenon and qualitatively reproduces the same patterns as in the Doodle data.

Standard theories in voting, and of Approval voting in particular, were mostly developed in the context of strategic voters who are trying to affect the identity of the selected candidate or alternative. Our analysis demonstrates that when voting in social settings, such as scheduling of a group meeting, voters may have other incentives that are just as important. Further understanding these social incentives and voter behavior, through empirical or theoretical analysis, is an important agenda of future research. It may help us to design better mechanisms for preference aggregation in everyday group decision making.

References

- [1] S Nageeb Ali and Navin Kartik. Herding with collective preferences. *Economic Theory*, 51(3):601–626, 2012.
- [2] Marco Battaglini. Sequential voting with abstention. *Games and Economic Behavior*, 51(2):445–463, 2005.
- [3] Steven J Brams and Peter C Fishburn. Approval voting. *The American Political Science Review*, pages 831–847, 1978.
- [4] William A Brock and Steven N Durlauf. Discrete choice with social interactions. *The Review of Economic Studies*, 68(2):235–260, 2001.
- [5] Steven Callander. Bandwagons and momentum in sequential voting. *The Review of Economic Studies*, 74(3):653–684, 2007.
- [6] Ioannis Caragiannis, Dimitris Kalaitzis, and Evangelos Markakis. Approximation algorithms and mechanism design for minimax approval voting. In *Proceedings of the National Conference on Artificial Intelligence*, 2010.

- [7] John R Chamberlin, Jerry L Cohen, and Clyde H Coombs. Social choice observed: Five presidential elections of the american psychological association. *The Journal of Politics*, 46(02):479–502, 1984.
- [8] Stephen Coleman. The effect of social conformity on collective voting behavior. *Political analysis*, 12(1):76–96, 2004.
- [9] Yvo Desmedt and Edith Elkind. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 347–356. ACM, 2010.
- [10] Ulle Endriss, Maria Silvia Pini, Francesca Rossi, and Kristen Brent Venable. Preference aggregation over restricted ballot languages: Sincerity and strategy-proofness. In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 122–127, 2009.
- [11] Jean-François Laslier. The leader rule: A model of strategic approval voting in a large electorate. *Journal of Theoretical Politics*, 21(1):113–136, 2009.
- [12] Jure Leskovec, Daniel P Huttenlocher, and Jon M Kleinberg. Governance in social media: A case study of the wikipedia promotion process. In *Proceedings of International Conference on Weblogs and Social Media*, 2010.
- [13] Ji Li and Lung-fei Lee. Binary choice under social interactions: an empirical study with and without subjective data on expectations. *Journal of Applied Econometrics*, 24(2):257–281, 2009.
- [14] Charles F Manski. Identification of endogenous social effects: The reflection problem. *The review of economic studies*, 60(3):531–542, 1993.
- [15] Nicholas Mattei, James Forshee, and Judy Goldsmith. An empirical study of voting rules and manipulation with large datasets. In *Fourth International Workshop on Computational Social Choice (COMSOC 2012)*. Springer, 2012.
- [16] Roger B. Myerson and Robert J. Weber. A theory of voting equilibria. *The American Political Science Review*, 87(1):102–114, 1993.
- [17] Michel Regenwetter and Bernard Grofman. Approval voting, borda winners, and condorcet winners: Evidence from seven elections. *Management Science*, 44(4):520–533, 1998.
- [18] Michel Regenwetter and Ilia Tsetlin. Approval voting and positional voting methods: Inference, relationship, examples. *Social Choice and Welfare*, 22(3):539–566, 2004.
- [19] Katharina Reinecke, Minh Khoa Nguyen, Abraham Bernstein, Michael Näf, and Krzysztof Z Gajos. Doodle around the world: Online scheduling behavior reflects cultural differences in time perception and group decision-making. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 45–54. ACM, 2013.
- [20] Ildikó Schlotter, Piotr Faliszewski, and Edith Elkind. Campaign management under approval-driven voting rules. In *Proceedings of the National Conference on Artificial Intelligence*, volume 11, pages 726–731, 2011.
- [21] Ad MA van Deemen and Noël P Vergunst. Empirical evidence of paradoxes of voting in dutch elections. *Public Choice*, pages 475–490, 1998.
- [22] Toby Walsh. Uncertainty in preference elicitation and aggregation. In *Proceedings of the National Conference on Artificial Intelligence*, volume 22, page 3. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press, 2007.

- [23] Robert J Weber. Approval voting. *The Journal of Economic Perspectives*, 9(1):39–49, 1995.
- [24] Fang Wu and Bernardo A Huberman. How public opinion forms. In *Internet and Network Economics*, pages 334–341. Springer, 2008.

James Zou
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA
Email: jzou@fas.harvard.edu

Reshef Meir
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA
Email: rmeir@seas.harvard.edu

David Parkes
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA
Email: parkes@eecs.harvard.edu