

Strategic Voting Behavior in Doodle Polls

James Zou
Microsoft Research
Cambridge, MA
jazo@microsoft.com

Reshef Meir
Harvard University
Cambridge, MA
rmeir@seas.harvard.edu

David C. Parkes
Harvard University
Cambridge, MA
parkes@eecs.harvard.edu

ABSTRACT

Finding a common time slot for a group event is a daily conundrum and illustrates key features of group decision-making. It is a complex interplay of individual incentives and group dynamics. A participant would like the final time to be convenient for her, but she is also expected to be cooperative towards other people's preferences. We combine large-scale data analysis with theoretical models from the voting literature to investigate strategic behaviors in event scheduling. We analyze all Doodle polls created in the US from July-September 2011 (over 340,000 polls), consisting of both hidden polls (a user cannot see other responses) and open polls (a user can see all previous responses). By analyzing the differences in behavior in hidden and open polls, we gain unique insights into strategies that people apply in a natural decision-making setting. 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.

Author Keywords

Event scheduling; strategic voting; Doodle; group dynamics.

ACM Classification Keywords

H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

INTRODUCTION

Online platforms for scheduling events are used more than ever. A recent survey has found that over 30% of Internet users use specialized online tools to schedule their meetings.¹ Among those tools, perhaps the most popular is the online calendar and scheduling website *Doodle* (www.doodle.com). In a typical Doodle poll, a *poll initiator* defines the possible dates and time slots (say, for an upcoming event), and sends a link to the poll to other participants, by e-mail or other means.

¹<http://en.blog.doodle.com/2011/07/13/>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CSCW 2015, March 14–18, 2015, Vancouver, BC, Canada.
Copyright © 2015 ACM 978-1-4503-2922-4/15/03 ...\$15.00.
<http://dx.doi.org/10.1145/2675133.2675273>

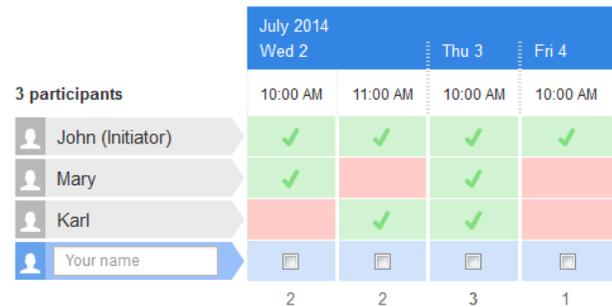


Figure 1: An open Doodle poll, as seen by the fourth voter.

Upon receiving the link, each participant can mark any number of time slots as available, see Figure 1. The poll initiator then uses the reported availability to determine the event time.

Much research has been dedicated to the adoption of scheduling tools, the conditions under which people are willing to use them, and how scheduling activities fit within the larger context of a workflow or a relationship [22, 25, 28]. We focus on the complementary question of user behavior during the scheduling process. Given that a group of people already decided to use a scheduling tool and have preferences over time slots, how do they report their availability? For example, we are interested in whether people apply strategic considerations, for example choosing to report as non-available some slots in which they could actually attend (but prefer not to).

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. Each participant may approve a subset of the slots by marking these as available. By doing so, a participant discloses partial information about her preferences, which is then aggregated with the votes of others. If we assume that the option approved by most participants gets selected, then this voting mechanism is known in the social choice literature as *Approval voting* [14] (where a voter can vote for as many alternatives as she likes). Strategic voting has been widely studied and discussed within the social choice literature, and so by rephrasing the group scheduling problem as a voting problem, we hope to gain important insights on the strategies that participants may apply.

In the simplest case, a participant may vote for exactly the time slots for which she is available. But in reality there is a

range of differing time slots, some of which may be feasible but inconvenient. How would she vote in this case? She 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 the meeting. Perhaps this is done in order to be cooperative, or because she personally benefits from a time that allows as many to attend as possible. Such sophisticated behavior can depend on how earlier participants have voted. Models from the voting literature offer predictions on how individuals would behave under these settings.

Doodle polls provide a unique opportunity to study the range of strategies that people apply in group scheduling situations. Consider the following three reasons. First, there is a huge database of polls that can be used for analysis. In this paper we use more than 14 million votes from 2 million responders in over 340,000 polls.

Second, the data 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 data as a baseline, we can study how voting behavior is affected by the group information available to the participant and avoid other confounds (for example, changes in responses that occur due to whether a response is placed early or late, and thus closer to the associated event). This also allows us to explore the implications of design decisions made by the initiator, in terms of setting a poll to be open versus closed.

Third, Approval voting has been extensively studied in the economics, political science 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.

Summary of Key Findings and Contributions

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 very unpopular time slots compared to hidden polls. Intermediate time slots have similar response rates between open and hidden polls.

We show that some of these empirical observations are inconsistent with standard models of Approval voting. We propose a new theoretical model, whereby a responder in an open poll votes for her preferred time slots, while also voting for highly popular slots that she does not disprefer, while also trying to appear cooperative, to explain the data.

Related work

Reinecke et al. [27] study how cultural differences between countries are reflected in scheduling patterns on Doodle. In particular they show that participants in collectivist countries tend to coordinate more with one another in open polls. Recognizing this, we use data from a single country (US) to avoid confounding due to such cultural differences.

The phenomenon of vote coordination in open online polls is related to a theoretical model of *herding* [5, 30, 17, 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 alternatives), and is learning about the best option from the responses of others rather than bringing her own preferences to bear.

The implications of sequential voting in more complex scenarios have also been studied [2, 9]. However these models deal with equilibrium analysis, thereby assuming very sophisticated participants. Moreover, these models do not consider Approval voting.

The role of incentives in scheduling has also received attention in the AI literature. However these papers focus mainly on the design of novel scheduling systems with various guarantees [10, 8]; or on optimizing the behavior of automated agents who operate on behalf of users [11, 7]. We are interested how human participants behave in an existing scheduling system.

Approval voting

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 all alternatives on which she votes over all other alternatives. Crucially, a voter may have more than one sincere way to vote. Brams and Fishburn 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 a voter has when deciding on her vote. Brams and Fishburn assume that voters have no knowledge whatsoever about the preferences or actions of others. A more elaborate model for Approval voting was studied by Weber [29], based on a general voting theory [23]. 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 sincere.

While the Weber model assumes that voters are highly sophisticated, capable of probabilistic calculations, a heuristic strategy called the *leader rule* has been suggested by Laslier [13].

²That is, alternatives can be partitioned to three sets, such the voter prefers all alternative in a set to all alternatives in the next set, but is indifferent between the alternatives in each set.

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

Other researchers have studied social factors affecting voting behavior, albeit not under the Approval system [6, 18]. In particular, the MBD model, named after Manski [19], Brock and Durlauf [4], assumes that a voter tries to meet the expectations of her peers (see [18]). The MBD model completely ignores any self-interested aspect of voting: the alternative that actually gets selected does not factor into the utility of the voter (assuming that the number of voters is high and thus every single vote is negligible). The simple model that we propose and support with the Doodle poll data combines both social and self-interest considerations.

EMPIRICAL RESULTS

Data collection

The dataset consists of all the polls created by US users on Doodle over the three month period July-September, 2011. It was collected and anonymized by Doodle and shared with the authors under a research agreement. We focus our analysis on polls with at least three participants, at least four time slots and only yes/no options. In Doodle, it is also possible to set up a poll where the participants can choose from one of three options (yes/no/maybe) for each time slot. However, there are very few three-option polls and we exclude them from our analysis. There are two types of 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 have 345,297 open polls and 7,390 hidden polls that pass this filtering. While there are fewer hidden polls, there are sufficiently many hidden polls to perform meaningful statistical analysis.

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. 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, where there is sufficient sample size for analysis.

Hypotheses

Our goal is to investigate how responders' behavior changes when all previous 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.*

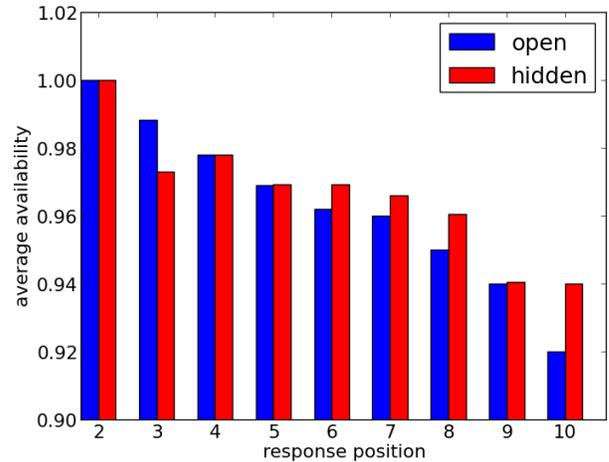


Figure 2: Normalized average availability as a function of responder position in hidden and open polls.

We hypothesize that Doodle responders, on average, will select time slots that work for other people, and will do this to a greater extent in open polls:

[H 1.2]. *The correlation of the votes of a participant and the aggregated votes from previous participants 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 (i.e. time slots that have received relatively large number of votes from previous participants), and will be more reluctant to approve unpopular alternatives:

[H 1.3]. *The probability that a voter will approve a popular alternative is higher in open polls than in hidden polls.*

[H 1.4]. *The probability that a voter will approve an unpopular alternative is lower in open polls than in hidden polls.*

Average availability

We define the *aggregate availability* of a poll to be the average proportion of ‘yes’ votes from all responders. 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 a linear regression, we still obtain significantly higher availability in open polls. In addition, we compute the average availability for each time slot, aggregated over all polls. We observe uniform availability across the time slots in both open and hidden polls, showing no bias in approvals for ‘leftmost slots’ (early times) or ‘rightmost slots’ (late times) in. 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 she approves. 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. It is also possible that busier people tend to respond later to polls and are also more constrained in the number of time slots they approve.

Both open and hidden polls have similar rate of decline. We see this in Figure 2. 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,³ and normalize availability by dividing the availability by responder 2’s availability.

Correlation with previous responders

For each responder we compute the Pearson correlation between her availability and the aggregated availability of all previous responses. This measures how likely she is to agree with earlier participants.

In both open and hidden polls, the correlation increases monotonically in response position (see Figure 3). Open polls have significantly higher correlation than hidden polls, supporting Hypothesis 1.2 ($p < 0.01$). This suggests that later participants actively try to match the previous responses that they observe. This finding is also consistent with Reinecke et al. [27], who found that participants in open polls tend to reach a consensus more often.

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 because 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.

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.

Which time slots get 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 the 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 4. The dots show the probability that the 11th responder inputs ‘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 relative to hidden polls for time slots with both low availability and high availability

³Indeed, responder 1 shows 12% and 20% higher availability than the poll average in hidden and open polls, respectively.

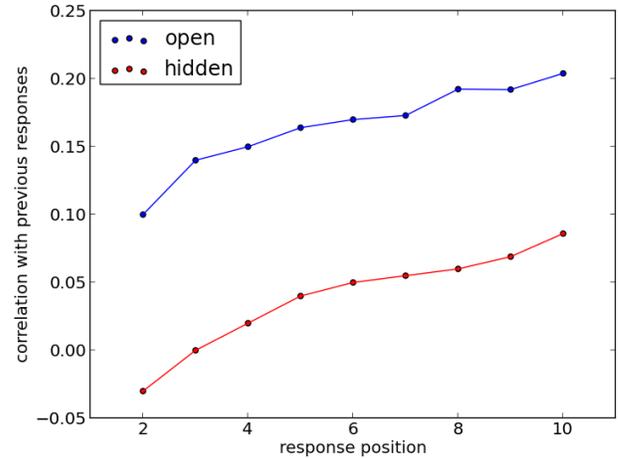


Figure 3: Correlation with previous responses in hidden and open polls.

($p < 0.01$). Both kinds of 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 5). There are 160,574 open polls and 5105 hidden polls with at least 6 participants. 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.

Conditional response rate

The response curves in Figures 4 and 5 are aggregated over all voters. Two possible explanations for 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 (possibly a minority group whose preferences differ from the rest). We compute conditional response rates and show that this supports 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 ($p < 0.01$).
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 ($p < 0.01$).

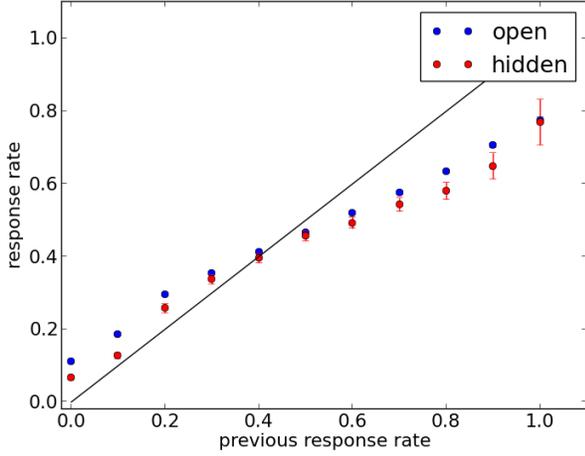


Figure 4: **Response curve for the 11th responder.** Each dot represents the probability that the 11th responder approves a slot (y-axis) if the slot has been approved by i out of the first 10 responders (x-axis), for $i = 0, \dots, 10$. Error bars shows the 95% confidence interval. The open poll error bars are small.

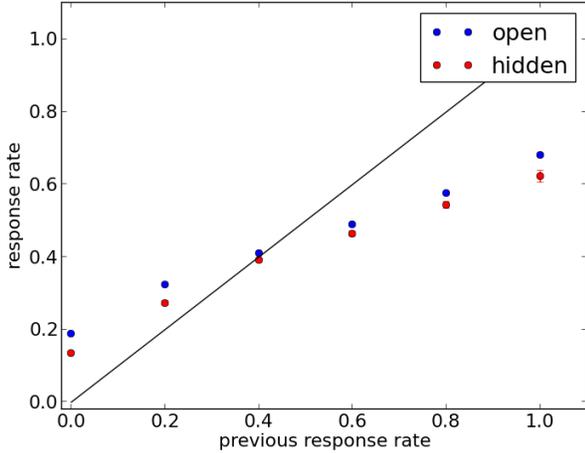


Figure 5: **Response curve for the 6th responder.** Each dot represents the probability that the 6th responder approves a slot (y-axis) if the slot has been approved by i out of the first 5 responders (x-axis), for $i = 0, \dots, 5$. In both Figures 4 and 5, open polls show statistically higher response rates for popular and unpopular time slots compared to hidden polls. For intermediate time slots, the difference between open and hidden polls is not statistically significant.

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.

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

ANALYSIS

The patterns observed in the previous section suggest that in open polls responders may resort to *strategic* behavior, whereby they do not simply mark all feasible time slots. In order to analyze possible incentives for such behavior, we introduce a preference model and concepts commonly used in the theory of Approval voting. This formal notation is required for a precise presentation and analysis, however we also explain each component in words.

Let $A = \{a_1, a_2, \dots, a_M\}$ denote the time slots designated by the initiator of the poll. 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 , denoted as r_n , is the set of slots that she approves. 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 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 $s(R) = (s(a, R))_{a \in A}$ the score vector over R . 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 $s_{\leq n}$ as a shorthand for $s(a, R_{\leq n})$ and $s(R_{\leq n})$, respectively. Given the first n voters, we divide the time slots with the highest, middle, and lowest number of votes into three sets $\{\text{Popular}_{\leq n}, \text{Intermediate}_{\leq n}, \text{Unpopular}_{\leq n}\}$.

The preference of voter v_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 $A_{k'}^n$, if k is less than k' . We allow the possibility that a set is empty.

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. 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 she prefers a to a' . In words: if all of the approved slots are preferred by the voter to all of the slots that are not approved.

A vote r'_n dominates vote r_n if for all actions of other voters, voter v_n prefers the outcome had she voted r'_n to the outcome had she voted r_n .

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 approves all top-level choices and no bottom-level choices. Given this 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 always 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.

Testing response functions

We describe some response functions that have been suggested in the literature.

In order to test some of the models (i.e., response functions), we run *social simulations* of open polls, albeit very simple ones. We assume that in hidden polls, responders submit admissible votes consistent with Prop. 1. Recall that by our other key assumption, the actual availability/preferences of participants in hidden and in open polls follow the same underlying distribution. This means we can use the open polls to derive some information on the preferences of participants in the hidden polls. In particular, to test a particular response function, we generate synthetic preference profiles using data from the hidden polls, and then run simulations of open polls where voters vote sequentially according to a particular response function and for each of the generated preference profiles. Then, we compare the patterns observed in the simulation with those described in the previous sections.

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 [26]. 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 $s_{\leq n-1}$, i.e. the votes placed so far. 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 [13] 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 preference profiles as follows. 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 with the same number of voters and slots by adopting the Leader Rule. Figure 7 (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 4, 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.

A SOCIAL VOTING MODEL

In many scenarios there are factors that affect the utility of the voter other than just the outcome. For example, voters may prefer an outcome that is also convenient for others [15], or vote in accordance to the expectations of their peers, as in the MBD model [19, 4]. Other considerations may bias

voters towards a strategy that is truthful or sincere, that is, in accordance with the voter’s own preferences [21, 24].

Social utility is one possible explanation for approving popular slots in a meeting context, as participants may prefer a time that is convenient to more people or allows more people to attend. But why would a responder mark an unpopular slot? Other than being able to see the responses of the previous participants, responders in open Doodle polls can also see their peer names, including after the poll is over. Social theories suggest when workers in a group monitor one another, they incentivize members of the group to hold up to the group norms [12].

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. In other words, bearing in mind that other participants can see her name and vote, a participant may want to approve more time slots to *appear more cooperative* and even if they are less convenient for herself. While any additional slot that is approved increases the social utility, voting on inconvenient slots is also risky for a participant, 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, $K = 3$, in defining Social voting:

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

In words, a voter approves all of her most preferred slots, irrespective of their popularity. In addition, among the slots at her second preference level, the voter approves those slots that are either very popular or very unpopular. No slot of the third preference level is approved.

By this theory, the Unpopular slots are marked to appear cooperative with the social norm, and since they have a low chance of getting selected anyways.⁴

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

The proof follows immediately from Proposition 1 of Brams and Fishburn [3]. The proposition shows that by following this Social voting response function, a voter can gain additional social utility without a substantial loss of “strategic” utility, and still remain sincere.

⁴A variation of the Social voting model is for the responder to also approve $A_3^n \cap \text{Unpopular}$ (note that this variation is neither admissible nor sincere). 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.

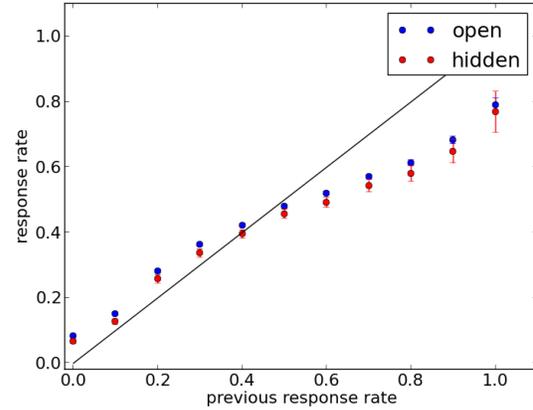


Figure 6: Response curve for synthetic open polls generated by the Leader rule. Hidden polls are from the real data.

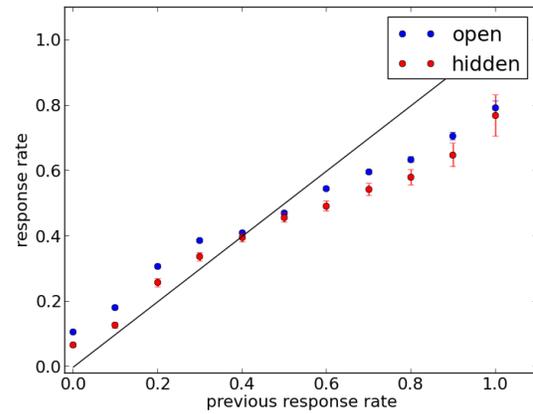


Figure 7: Response curve for synthetic open polls generated by the Social voting rule (compare with the Leader rule in Figure 6). The Leader rule does not generate data matching the observed response rates, while the Social voting rule does reproduce the observed pattern.

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 (all most preferred slots). 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 when participants adopt the Social voting rule 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 Popular _{$\leq n-1$} , Intermediate _{$\leq n-1$} and

Unpopular $_{\leq n-1}$ and approves $\hat{A}_1^n \cup (\hat{A}_2^n \cap (\text{Popular}_{\leq n-1} \cup \text{Unpopular}_{\leq n-1}))$, as specified by the Social voting response function.

Figures 6 and 7 show the $n = 11$ response curve for these synthetic open polls under this behavior model, combined with the curve for real hidden polls. 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. The pattern we see in the social simulations matches the pattern observed in the real data (see Figure 4).

DISCUSSION

Standard models of user behavior in voting were mostly developed in the context of strategic voters who are trying to affect the identity of the selected candidate or alternative. However social factors beyond those that relate to the choice of the winning candidate can play some role in voters' decisions. Our analysis demonstrates that in social settings, such as scheduling a group meeting, voters may have other incentives that are also important.

We show 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 need to approve it, in order to be cooperative or because she personally benefits from a time that allows as many to attend as possible. 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. All of our analysis are based on data from the Doodle platform. It would be interesting to see if similar patterns can be found in other online and offline scheduling systems.

In future work we intend to further explore and model social utility. In particular, we would like to have a distinction between the social utility gained from *being* more cooperative to that gained from *appearing* as such. A complete behavioral model should take into account both the utility structure and the epistemic state of the voter [23, 20]. In that respect, it is important to mention that strategic voting in itself is not *a priori* harmful, and may often increase the overall social welfare of the group [16]. Further research is required to determine the welfare implications of various strategic behavior patterns in Doodle.

Another interesting direction is to understand more fine-grained social dynamics in online scheduling. For example, if one has observations of the same group of individuals over multiple polls, then it should be possible to infer the influence of each member in terms of the effects of her votes on other members' votes. It would then be interesting to investigate whether influential members tends to vote early or late, as well as how members try to maximize their influence.

Finally, further understanding these social incentives and

voter behavior, through empirical or theoretical analysis, also has important implications for the design of group interaction mechanisms. For example, Doodle currently offers a binary choice between hidden polls (where voters' responses are hidden both throughout the poll and after it) and open polls. Our results suggest that while open polls provide opportunity for coordination, they also provide incentives for strategic behavior that does not benefit the group and could lead to miscoordination (marking unpopular slots). A clear separation between these features and their effects may allow the design of better mechanisms for preference aggregation in everyday group decision making.

Acknowledgement We would like to David Gubler at Doodle for extracting the data and Doodle and Katharina Reinecke for sharing the data. We have also benefited from insightful discussions with Katharina Reinecke, Krzysztof Gajos and Myke Näf.

Parkes was supported in part by NSF grant CCF #1301976 and the SEAS TomKat fund.

REFERENCES

1. Ali, S. N., and Kartik, N. Herding with collective preferences. *Economic Theory* 51, 3 (2012), 601–626.
2. Battaglini, M. Sequential voting with abstention. *Games and Economic Behavior* 51, 2 (2005), 445–463.
3. Brams, S. J., and Fishburn, P. C. Approval voting. *The American Political Science Review* (1978), 831–847.
4. Brock, W. A., and Durlauf, S. N. Discrete choice with social interactions. *The Review of Economic Studies* 68, 2 (2001), 235–260.
5. Callander, S. Bandwagons and momentum in sequential voting. *The Review of Economic Studies* 74, 3 (2007), 653–684.
6. Coleman, S. The effect of social conformity on collective voting behavior. *Political analysis* 12, 1 (2004), 76–96.
7. Crawford, E., and Veloso, M. Opportunities for learning in multi-agent meeting scheduling. In *Proceedings of the AAAI 2004 Symposium on Artificial Multiagent Learning, Washington, DC* (2004).
8. Crawford, E., and Veloso, M. Mechanism design for multi-agent meeting scheduling. *Web Intelligence and Agent Systems* 4, 2 (2006), 209–220.
9. Desmedt, Y., and Elkind, E. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM conference on Electronic commerce*, ACM (2010), 347–356.
10. Ephrati, E., Zlotkin, G., and Rosenschein, J. S. A non-manipulable meeting scheduling system. In *Proceedings of the 13th international workshop on distributed artificial intelligence* (1994), 105–125.
11. Garrido, L., and Sycara, K. Multi-agent meeting scheduling: Preliminary experimental results. In

- Proceedings of the Second International Conference on Multiagent Systems* (1996), 95–102.
12. Kandel, E., and Lazear, E. P. Peer pressure and partnerships. *Journal of political Economy* 100, 4 (1992), 801.
 13. Laslier, J.-F. The leader rule: A model of strategic approval voting in a large electorate. *Journal of Theoretical Politics* 21, 1 (2009), 113–136.
 14. Laslier, J.-F., and Sanver, M. R. *Handbook on approval voting*. Springer, 2010.
 15. Lee, H. Algorithmic and game-theoretic approaches to group scheduling. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, International Foundation for Autonomous Agents and Multiagent Systems (2014), 1709–1710.
 16. Lehtinen, A. The welfare consequences of strategic behaviour under approval and plurality voting. *European Journal of Political Economy* 24, 3 (2008), 688–704.
 17. Leskovec, J., Huttenlocher, D. P., and Kleinberg, J. M. Governance in social media: A case study of the wikipedia promotion process. In *Proceedings of International Conference on Weblogs and Social Media* (2010).
 18. Li, J., and Lee, L.-f. Binary choice under social interactions: an empirical study with and without subjective data on expectations. *Journal of Applied Econometrics* 24, 2 (2009), 257–281.
 19. Manski, C. F. Identification of endogenous social effects: The reflection problem. *The review of economic studies* 60, 3 (1993), 531–542.
 20. Meir, R., Lev, O., and Rosenschein, J. S. A local-dominance theory of voting equilibria. In *Proceedings of the fifteenth ACM conference on Economics and computation*, ACM (2014), 313–330.
 21. Meir, R., Polukarov, M., Rosenschein, J. S., and Jennings, N. R. Convergence to equilibria in plurality voting. In *AAAI'10* (2010).
 22. Mosier, J. N., and Tammara, S. G. When are group scheduling tools useful? *Computer Supported Cooperative Work (CSCW)* 6, 1 (1997), 53–70.
 23. Myerson, R. B., and Weber, R. J. A theory of voting equilibria. *The American Political Science Review* 87, 1 (1993), 102–114.
 24. Obratsova, S., Markakis, E., and Thompson, D. R. Plurality voting with truth-biased agents. In *SAGT'13*. 2013, 26–37.
 25. Palen, L. Social, individual and technological issues for groupware calendar systems. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, ACM (1999), 17–24.
 26. Regenwetter, M., and Tsetlin, I. Approval voting and positional voting methods: Inference, relationship, examples. *Social Choice and Welfare* 22, 3 (2004), 539–566.
 27. Reinecke, K., Nguyen, M. K., Bernstein, A., Näf, M., and Gajos, K. Z. 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*, ACM (2013), 45–54.
 28. Thayer, A., Bietz, M. J., Derthick, K., and Lee, C. P. I love you, let's share calendars: calendar sharing as relationship work. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, ACM (2012), 749–758.
 29. Weber, R. J. Approval voting. *The Journal of Economic Perspectives* 9, 1 (1995), 39–49.
 30. Wu, F., and Huberman, B. A. How public opinion forms. In *Internet and Network Economics*. Springer, 2008, 334–341.