

Where are the really hard manipulation problems? The manipulation phase transition*

Toby Walsh

NICTA and UNSW

Sydney, Australia

toby.walsh@nicta.com.au

Abstract

Voting is a simple mechanism to aggregate the preferences of agents. Many voting rules have been shown to be NP-hard to manipulate. However, a number of recent theoretical results have suggested that this complexity may only be in the worst-case and manipulation may be easy in practice. In this paper, we show that empirical studies are useful in improving our understanding of this issue. We demonstrate that there is a smooth transition in the probability that a coalition can elect a desired candidate as the size of the manipulating coalition is varied. We show that a rescaled probability curve displays a simple and universal form independent of the size of the problem. We argue that for many independent and identically distributed votes, manipulation will be computationally easy even when the coalition of manipulators is critical in size. Based on this argument, we identify a situation in which manipulation is computationally hard. This is when votes are highly correlated and the election is “hung”. We show, however, that even a single uncorrelated voter is enough to make manipulation easy again.

1 Introduction

The Gibbard-Satterthwaite theorem proves that, under some simple assumptions, most voting rules are manipulable. That is, it may pay for an agent not to report their preferences truthfully. One possible escape from this result was proposed by Bartholdi, Tovey and Trick [Bartholdi *et al.*, 1989]. Whilst a manipulation may exist, perhaps it is computationally too difficult to find? Many results have subsequently been proven showing that various voting rules are NP-hard to manipulate under different assumptions including: an unbounded number of candidates; a small number of candidates but weighted votes; and uncertainty in the distribution of votes. See, for instance, [Bartholdi *et al.*, 1989; Bartholdi and Orlin, 1991;

*NICTA is funded by the Australian Government through the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

Conitzer *et al.*, 2007]. There is, however, increasing concern that worst-case results like these may not reflect the difficulty of manipulation in practice. Indeed, a number of recent theoretical results suggest that manipulation may often be computationally easy [Conitzer and Sandholm, 2006; Procaccia and Rosenschein, 2007b; Xia and Conitzer, 2008a; Friedgut *et al.*, 2008; Xia and Conitzer, 2008b].

In this paper we show that, in addition to attacking this question theoretically, we can profitably study it empirically. There are several reasons why empirical analysis is useful. First, theoretical analysis is often asymptotic and so does not show the size of any hidden constants. In addition, elections are typically bounded in size. Can we be sure that asymptotic behaviour is relevant for the finite sized electorates met in practice? Second, theoretical analysis is often restricted to particular distributions (e.g. independent and identically distributed votes). Manipulation may be very different in practice due to correlations between the votes. For instance, if all preferences are single-peaked then there are voting rules which cannot be manipulated. With such rules, it is in the best interests of all agents to state their true preferences. Third, many of these theoretical results about the easiness of manipulation have been hard won and are limited in their applicability. For instance, Friedgut *et al.* have not so far been able to extend their result beyond three candidates [Friedgut *et al.*, 2008]. An empirical study may quickly suggest if the result extends to more candidates. Finally, empirical studies may suggest new avenues for theoretical study. For example, the experiments reported here suggest a simple and universal form for the probability that a coalition of critical size is able to elect a desired candidate. It would be interesting to try to tackle this theoretically.

2 Background

We suppose that there are n agents who have voted and a coalition of m additional agents who wish to manipulate the result. When the manipulating coalition is small, they have too little weight to be able to change the result. On the other hand, when the coalition is large, they are sure to be able to make their desired candidate win. Procaccia and Rosenschein prove that for most scoring rules and a wide variety of distributions over votes, when $m = o(\sqrt{n})$, the probability that a manipulating coalition can change the result tends to 0, and when $m = \omega(\sqrt{n})$, the probability that they can manipu-

late the result tends to 1 [Procaccia and Rosenschein, 2007a]. They offer two interpretations of this result. On the positive side, they suggest it may focus attention on other distributions which are computationally hard to manipulate. On the negative side, they suggest that it may strengthen the argument that manipulation problems are easy on average.

More recently, Xia and Conitzer show that for a large class of voting rules, as the number of agents grows, either the probability that a coalition can manipulate the result is very small (as the coalition is too small), or the probability that they can (easily) manipulate the result to make any alternative win is very large [Xia and Conitzer, 2008a]. They leave open only a small interval in the size of the coalition for which the coalition is large enough to be able to manipulate but not obviously large enough to be able to manipulate the result easily. More precisely, for a wide range of voting rules including scoring rules, STV, Copeland and maximin, with votes which are drawn independently and with an identical distribution that is positive everywhere, they identify three cases:

- if $m = O(n^p)$ for $p < \frac{1}{2}$ then the probability that the result can be changed is $O(\frac{1}{\sqrt{n}})$;
- if $m = \Omega(n^p)$ for $p > \frac{1}{2}$ and $o(n)$ and votes are uniform then the probability that the result can be manipulated is $1 - O(e^{-\Theta(n^{2p-1})})$ using a simple greedy procedure;
- if $m = \Theta(\sqrt{n})$ then they provide no result.

In this paper, we shall provide empirical evidence to help close this gap and understand what happens when the coalition is of a critical size that grows as $\Theta(\sqrt{n})$.

3 Finding manipulations

We will focus on the veto rule. This is a scoring rule in which each agent gets to cast a veto against one candidate. The candidate with the fewest vetoes wins. We suppose that tie-breaking is in favor of the manipulators. However, it is easy to relax this assumption. There are several reasons why we start this investigation into the complexity of manipulation with the veto rule. First, it is very simple to reason about. This can be contrasted with other voting rules that are computationally hard to manipulate. For example, the STV rule is NP-hard to manipulate [Bartholdi and Orlin, 1991] but its complexity appears to come in part from reasoning about what happens between the different rounds. Second, it is on the borderline of complexity since constructive manipulation of the veto rule by a coalition of weighted agents is NP-hard but destructive manipulation is polynomial [Conitzer *et al.*, 2007]. Third, efficient number partitioning algorithms can be used to compute a successful manipulation. In particular, we show that manipulation of an election with 3 candidates and weighted votes (which is NP-hard [Conitzer *et al.*, 2007]) can be very directly reduced to 2-way number partitioning (which we can solve using the efficient CKK algorithm [Korf, 1995]). A similar argument can be given to show that the manipulation of a veto election of p candidates by a weighted coalition can be reduced to finding a $p - 1$ -way partition of numbers.

Theorem 1 *There exists a successful manipulation of an election with 3 candidates by a weighted coalition using the*

veto rule iff there exists a partitioning of $W \cup \{|a - b|\}$ into two bags such that the difference between their two sums is less than or equal to $a + b - 2c + \sum_{i \in W} i$ where W is the multiset of weights of the manipulating coalition, a , b and c are the weights of vetoes assigned to the three candidates by the non-manipulators and the manipulators wish the candidate with weight c to win.

Proof: It never helps a coalition manipulating the veto rule to veto the candidate that they wish to win. The coalition does, however, need to decide how to divide their vetoes between the candidates that they wish to lose. Consider the case $a \geq b$. Suppose the partition has weights $w - \Delta/2$ and $w + \Delta/2$ where $2w = \sum_{i \in W \cup \{|a-b|\}} i$ and Δ is the difference between the two sums. The same partition of vetoes is a successful manipulation iff the winning candidate has no more vetoes than the next best candidate. That is, $c \leq b + (w - \Delta/2)$. Hence $\Delta \leq 2w + 2b - 2c = (a - b) + 2b - 2c + \sum_{i \in W} i = (a + b - 2c) + 2 \sum_{i \in W} i$. In the other case, $a < b$ and $\Delta \leq (b + a - 2c) + \sum_{i \in W} i$. Thus $\Delta \leq a + b - 2c + \sum_{i \in W} i$. \square

A similar (but slightly more complex) argument can be used to show that manipulation of *any* scoring rule with 3 candidates and weighted votes can be reduced to 2-way number partitioning. However, the argument used does not extend to more than 3 candidates. Manipulating elections with greater than 3 candidates and scoring rules other than veto or plurality seems to require other more complex methods.

4 Uniform votes

We consider the case that the n agents veto uniformly at random one of the 3 possible candidates, and vetoes carry weights drawn uniformly from $(0, k]$. When the coalition is small in size, it has too little weight to be able to change the result. On the other hand, when the coalition is large in size, it is sure to be able to make a favored candidate win. There is thus a transition in the manipulability of the problem as the coalition size increases (see Figure 1).

Based on [Procaccia and Rosenschein, 2007a; Xia and Conitzer, 2008a], we expect the critical coalition size to increase as \sqrt{n} . In Figure 2, we see that the phase transition appears to display a simple and universal form when plotted against m/\sqrt{n} . The phase transition appears to be smooth, with the probability varying slowly and not approaching a step function as problem size increases. We obtained a good fit with $1 - 2/3e^{m/\sqrt{n}}$. Other smooth phase transitions have been seen with 2-coloring [Achlioptas, 1999], 1-in-2 satisfiability and Not-All-Equal 2-satisfiability [Walsh, 2002]. It is interesting to note that all these decision problems are polynomial.

The theoretical results mentioned earlier leave open how hard it is to compute whether a manipulation is possible when the coalition size is critical. Figure 3 displays the computational cost to find a manipulation (or prove none exists) using the efficient CKK algorithm. Even in the critical region where problems may or may not be manipulable, it is easy to compute whether the problem is manipulable. All problems can be solved in a few branches. This contrasts with phase transition behaviour in problems like propositional satisfiability

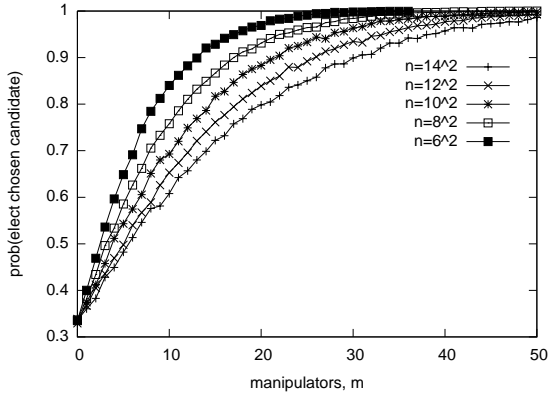


Figure 1: Probability of a coalition of m agents electing a chosen candidate where n agents have already voted, and vetoes are weighted and uniformly drawn from $(0, 2^8]$. Note that at $m = 0$, there is a 1/3rd chance that the non-manipulators have already elected this candidate. In this and all subsequent experiments, we tested 10,000 problems at each data point.

[Cheeseman *et al.*, 1991; Mitchell *et al.*, 1992] and number partitioning [Gent and Walsh, 1998], where the hardest problems tend to occur around the phase transition.

5 Why hard problems are rare

By using the reduction of manipulation problems to number partitioning, we give a heuristic argument why hard manipulation problems are vanishing rare as $n \rightsquigarrow \infty$ and $m = \Theta(\sqrt{n})$. The basic idea is simple: by the time the coalition is large enough to be able to change the result, the variance in scores between the candidates is likely to be so large that computing a successful manipulation or proving none is possible will be easy.

Suppose that the manipulators want candidates A and B to lose so that C wins, and that the non-manipulators have cast vetoes of weight a, b and c for A, B and C respectively. Without loss of generality we suppose that $a \geq b$. There are three cases to consider. In the first case, $a \geq c$ and $b \geq c$. It is then easy for the manipulators to make C win since C wins whether they veto A or B . In the second case, $a \geq c > b$. Again, it is easy for the manipulators to decide if they can make C win. They all veto B . There is a successful manipulation iff C now wins. In the third case, $a < c$ and $b < c$. The manipulators must partition their m vetoes between A and B so that the total vetoes received by A and B exceeds those for C . Let d be the deficit in weight between A and C and between B and C . That is, $d = (c-a) + (c-b) = 2c - a - b$. We can model d as the sum of n random variables drawn uniformly with probability 1/3 from $[0, 2k]$ and with probability 2/3 from $[-k, 0]$. These variables have mean 0 and variance $2k^2/3$. By the Central Limit Theorem, d tends to a normal distribution with mean 0, and variance $s^2 = 2nk^2/3$. For a manipulation to be possible, d must be less than w , the sum of the weights of the vetoes of the manipulators. By the Central Limit Theorem, w also tends to a normal distribution with




Figure 2: Rescaled probability that a coalition of m agents can elect a chosen candidate where n agents have already voted, and vetoes are weighted and uniformly drawn from $(0, 2^8]$. The x-axis is scaled by $1/\sqrt{n}$.




Figure 3: Computational cost for the CKK algorithm to decide if a coalition of m agents can manipulate a veto election where n agents have already voted, and vetoes are weighted and uniformly drawn from $(0, 2^m]$. Even the most difficult problems are solved with almost no search.

mean $\mu = mk/2$, and variance $\sigma^2 = 2mk^2/3$.

A simple heuristic argument due to [Karmarkar *et al.*, 1986] and also based on the Central Limit Theorem upper bounds the optimal partition difference of m numbers from $(0, k]$ by $O(k\sqrt{m}/2^m)$. In addition, based on the phase transition in number partitioning [Gent and Walsh, 1998], we expect partitioning problems to be easy unless $\log_2(k) = \Theta(m)$. Combining these two observations, we expect hard manipulation problems when $0 \leq w - d \leq \alpha\sqrt{m}$ for some constant α . The probability of this occurring is:

$$\int_0^\infty \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \int_{x-\alpha\sqrt{m}}^x \frac{1}{\sqrt{2\pi}s} e^{-\frac{y^2}{2s^2}} dy dx$$

By substituting for s, μ and σ , we get:

$$\int_0^\infty \frac{1}{\sqrt{4\pi mk^2/3}} e^{-\frac{(x-mk/2)^2}{4mk^2/3}} \int_{x-\alpha\sqrt{m}}^x \frac{1}{\sqrt{4\pi nk^2/3}} e^{-\frac{y^2}{4nk^2/3}} dy dx$$

For $n \rightsquigarrow \infty$, this tends to:

$$\int_0^\infty \frac{1}{\sqrt{4\pi mk^2/3}} e^{-\frac{(x-mk/2)^2}{4mk^2/3}} \frac{\alpha\sqrt{m}}{\sqrt{4\pi nk^2/3}} e^{-\frac{x^2}{4nk^2/3}} dx$$

As $e^{-z} \leq 1$ for $z > 0$, this is upper bounded by:

$$\frac{\alpha\sqrt{m}}{\sqrt{4\pi nk^2/3}} \int_0^\infty \frac{1}{\sqrt{4\pi mk^2/3}} e^{-\frac{(x-mk/2)^2}{4mk^2/3}} dx$$

Since the integral is bounded by 1, $m = \Theta(\sqrt{n})$ and $\log_2(k) = \Theta(m)$, this upper bound varies as:

$$O\left(\frac{1}{\sqrt{m}2^m}\right)$$

Thus, we expect hard instances of manipulation problems to be exponentially rare. Since even a brute force manipulation algorithm takes $O(2^m)$ time in the worst-case, we do not expect the hard instances to have a significant impact on the average-case as n (and thus m) grows. We stress this is only a heuristic argument. It makes assumptions about the complexity of manipulation problems (in particular that hard instances should lie within the narrow interval $0 \leq w - d \leq \alpha\sqrt{m}$). These assumptions are only supported by empirical observation and informal argument. However, the experimental results reported in Figure 3 support the overall conclusions.

6 Varying weights

The theoretical analyses of manipulation in [Procaccia and Rosenschein, 2007a; Xia and Conitzer, 2008a] suggest that the probability of an election being manipulable is largely independent of k , the size of the weights attached to the vetoes. Figure 4 demonstrates that this indeed appears to be the case in practice. When weights are varied in size from 2^8 to 2^{16} ,




Figure 4: Independence of the size of the weights and the manipulability of an election. Probability that a coalition of m agents can elect a chosen candidate where n agents have already voted, and vetoes are weighted and uniformly drawn from $(0, k]$.

the probability does not appear to change. In fact, the probability curve fits the same simple and universal form plotted in Figure 2. We also observed that the cost of computing a manipulation or proving that none is possible did not change as the weights were varied in size.

7 Normally distributed votes

What happens with other distributions of votes? The theoretical analyses of manipulation in [Procaccia and Rosenschein, 2007a; Xia and Conitzer, 2008a] suggest that there is a critical coalition size that increases as $\Theta(\sqrt{n})$ for many types of independent and identically distributed random votes. Similarly, our heuristic argument about why hard manipulation problems are vanishingly rare depends on application of the Central Limit Theorem. It therefore works with other types of independent and identically distributed random votes.




Figure 5: Weighted votes taken from a normal distribution. We plot the probability that a coalition of m agents can elect a chosen candidate where n agents have already voted, and vetoes are weighted and drawn from a normal distribution with mean 2^8 and standard deviation 2^7 . The x-axis is scaled by \sqrt{n} .

We shall consider therefore another type of independent and identically distributed vote. In particular, we study an election in which weights are independently drawn from a normal distribution. Figure 5 shows that there is again a smooth phase transition in manipulability. We also plotted Figure 5 on top of Figures 2 and 4. All curves appear to fit the same simple and universal form. As with uniform weights, the computational cost of deciding if an election is manipulable was small even when the coalition size was critical. Finally, we varied the parameters of the normal distribution. The probability of electing a chosen candidate as well as the cost of computing a manipulation did not appear to depend on the mean or variance of the distribution.

8 Correlated votes

To find hard manipulation problems, it seems we must look to votes which are more correlated. For example, consider a “hung” election where all n agents veto the candidate that the manipulators wish to win, but the m manipulators have exactly twice the weight of vetoes of the n agents. This election is finely balanced. The favored candidate of the manipulators wins iff the manipulators perfectly partition their vetoes between the two candidates that they wish to lose.

In Figure 6, we plot the probability that the m manipulators can make their preferred candidate win in such a “hung” election as we vary the size of their weights k . Similar to number




Figure 6: Manipulation of an election where votes are highly correlated and the result is “hung”. We plot the probability that a coalition of m agents can elect a chosen candidate where the vetoes of the manipulators are weighted and uniformly drawn from $(0, k]$, the other agents have all vetoed the candidate that the manipulators wish to win, and the sum of the weights of the manipulators is twice that of the non-manipulators.

partitioning [Gent and Walsh, 1998], we see a rapid transition in manipulability around $\log_2(k)/m \approx 1$. In Figure 7, we observe that there is a rapid increase in the computationally complexity to compute a manipulation around this point.

What happens when the votes are not so perfectly correlated? We consider an election which is perfectly hung as before except for one agent who votes at random between the three candidates. In Figure 8, we plot the cost of computing a manipulation as k^l , the size of the weight of this single random veto increases. We see that even one uncorrelated vote is enough to make manipulation easy if it has the same magnitude in weight as the vetoes of the manipulators. This suggests that we will only find hard manipulation problems in highly correlated voting distributions.

9 Other related work

There have been a number of other recent theoretical results about the computational complexity of manipulating elections. For instance, Procaccia and Rosenschein give a simple greedy procedure that will in polynomial time find a manipulation of a scoring rule for any “junta” distribution of weighted votes with a probability of failure that is an inverse polynomial in n [Procaccia and Rosenschein, 2007b]. A “junta” distribution is concentrated on the hard instances.

As a second example, Friedgut, Kalai and Nisan prove that if the voting rule is neutral and far from any dictatorship and there are three candidates then there exists an agent for whom a random manipulation succeeds with probability $\Omega(\frac{1}{n})$ where n is the number of agents [Friedgut *et al.*, 2008]. They were, however, unable to extend their proof to four (or more) candidates. Xia and Conitzer showed that, starting from different assumptions, a random manipulation would succeed with probability $\Omega(\frac{1}{n})$ for 3 or more candidates for STV, for 4 or more candidates for a scoring rule and for 5 or




Figure 7: The cost to decide if a hung election can be manipulated. We plot the computational cost for the CCK algorithm to decide if a coalition of m agents can manipulate a veto election where the vetoes of the manipulators are weighted and uniformly drawn from $(0, k]$, the other agents have all vetoed the candidate that the manipulators wish to win, and the sum of the weights of the manipulators is twice that of the non-manipulators.

more candidates for Copeland [Xia and Conitzer, 2008b].

Coleman and Teague provide polynomial algorithms to compute a manipulation for the STV rule when either the number of voters or the number of candidates is fixed [Coleman and Teague, 2007]. They also conducted an empirical study which demonstrates that only relatively small coalitions are needed to change the elimination order of the STV rule. They observe that most uniform and random elections are not trivially manipulable using a simple greedy heuristic.

Finally, a similar phenomena has been observed in the phase transition for deciding if a random graph contains a Hamiltonian cycle [Vandegriend and Culberson, 1998]. If the number of edges is small, there is likely to be a node of degree smaller than 2. There cannot therefore be any Hamiltonian cycle. By the time that there are enough edges for all nodes to be degree two, there are likely to be many possible Hamiltonian cycles and even a simple heuristic can find one. Thus, the phase transition in the existence of a Hamiltonian cycle is not associated with hard instances of the problem. We saw a similar phenomenon here. By the time the coalition is large enough to manipulate the result, the variance in scores between the candidates is likely to be so large that computing a successful manipulation or proving none is possible is easy.

10 Conclusions

We have studied the question of whether computational complexity is a barrier to the manipulation of a voting rule. We showed that there is a *smooth* transition in the probability that a coalition can elect a desired candidate as the size of the manipulating coalition is varied. We demonstrated that a rescaled probability curve displays a simple and universal form independent of the size of the problem. Unlike phase transitions for other NP-complete problems, hard problems are not associated with the transition between satisfiable and




Figure 8: The impact of a single random voter on the manipulability of a hung election. We plot the computational cost for the CKK algorithm to decide if a coalition of m agents can manipulate a veto election where the vetoes of the manipulators are weighted and uniformly drawn from $(0, k]$, the non-manipulating agents have all vetoed the candidate that the manipulators wish to win, and the sum of the weights of the manipulators is twice that of the non-manipulators except for one random non-manipulating agent whose weight is uniformly drawn from $(0, k']$. When the veto of the one random voter has the same weight as the other voters, it is computationally easy to decide if the election can be manipulated.

unsatisfiable problems. We observed similar behavior with other independent and identically distributed votes like those following a normal distribution. Finally, we studied the impact of correlation between votes. We showed that manipulation is computationally hard when votes are highly correlated and the election is “hung”. However, even a single uncorrelated voter was enough to make manipulation easy again.

What general lessons can be learnt from this study? First, whilst we have focused on the veto rule, it is likely that similar behavior will be seen with other voting rules. It would, for instance, be interesting to study the STV rule. This is NP-hard to manipulate even without weights. In addition, the rule has multiple rounds making it hard to reason about and to manipulate. Second, there appears to be an universal form for the probability that a coalition is able to elect a chosen candidate. It would be interesting to derive this form theoretically. Third, we conjecture that there is a connection between the smoothness of the phase transition and problem hardness. Sharp phase transitions (like that for propositional satisfiability) appear to be associated with hard instances of decision problems, whilst smooth transitions appear to be associated with easy instances of NP-hard problems or with polynomial problems like 2-colorability. Fourth, these results demonstrate that empirical studies can improve our understanding of manipulation. It would therefore be interesting to consider similar studies of related topics like bribery and control.

References

[Achlioptas, 1999] D. Achlioptas. *Threshold phenomena in random graph colouring and satisfiability*. PhD thesis,

Dept. of Computer Science, University of Toronto, 1999.

[Bartholdi and Orlin, 1991] J.J. Bartholdi and J.B. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

[Bartholdi et al., 1989] J.J. Bartholdi, C.A. Tovey, and M.A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[Cheeseman et al., 1991] P. Cheeseman, B. Kanefsky, and W.M. Taylor. Where the really hard problems are. In *Proc. of 12th IJCAI*, pages 331–337. 1991.

[Coleman and Teague, 2007] T. Coleman and V. Teague. On the complexity of manipulating elections. In *Proc. of 13th Australasian Theory Symposium*, 2007.

[Conitzer and Sandholm, 2006] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proc. of 21st Nat. Conf. on AI. AAAI*, 2006.

[Conitzer et al., 2007] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate. *JACM*, 54, 2007.

[Friedgut et al., 2008] E. Friedgut, G. Kalai, and N. Nisan. Elections can be manipulated often. In *Proc. 49th FOCS*. 2008.

[Gent and Walsh, 1998] I.P. Gent and T. Walsh. Analysis of heuristics for number partitioning. *Computational Intelligence*, 14(3):430–451, 1998.

[Karmarkar et al., 1986] N. Karmarkar, R. Karp, J. Lueker, and A. Odlyzko. Probabilistic analysis of optimum partitioning. *J. of Applied Probability*, 23:626–645, 1986.

[Korf, 1995] R. Korf. From approximate to optimal solutions: A case study of number partitioning. In *Proc. of 14th IJCAI*. 1995.

[Mitchell et al., 1992] D. Mitchell, B. Selman, and H. Levesque. Hard and Easy Distributions of SAT Problems. In *Proc. of 10th Nat. Conf. on AI. AAAI*, 1992.

[Procaccia and Rosenschein, 2007a] A. D. Procaccia and J. S. Rosenschein. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proc. of AAMAS-07*, 2007.

[Procaccia and Rosenschein, 2007b] A. D. Procaccia and J. S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *JAIR*, 28:157–181, 2007.

[Vandegriend and Culberson, 1998] B. Vandegriend and J. Culberson. The $G(n,m)$ phase transition is not hard for the Hamiltonian cycle problem. *JAIR*, 9:219–245, 1998.

[Walsh, 2002] T. Walsh. From P to NP: COL, XOR, NAE, 1-in-k, and Horn SAT. In *Proc. of 17th Nat. Conf. on AI. AAAI*, 2002.

[Xia and Conitzer, 2008a] Lirong Xia and Vincent Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proc. of 9th ACM Conf. on Electronic Commerce*, 2008.

[Xia and Conitzer, 2008b] Lirong Xia and Vincent Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proc. of 9th ACM Conf. on Electronic Commerce*, 2008.