



Electing the Doge of Venice: analysis of a 13th Century protocol[♦]

Miranda Mowbray, Dieter Gollmann¹
Enterprise Systems and Storage Laboratory
HP Laboratories Bristol
HPL-2007-28(R.1)
July 12, 2007*

voting theory,
leader election,
Venice

This paper discusses the protocol used for electing the Doge of Venice between 1268 and the end of the Republic in 1797. We will show that it has some useful properties that in addition to being interesting in themselves, also suggest that its fundamental design principle is worth investigating for application to leader election protocols in computer science. For example it gives some opportunities to minorities while ensuring that more popular candidates are more likely to win, and offers some resistance to corruption of voters.

The most obvious feature of this protocol is that it is complicated and would have taken a long time to carry out. We will advance a hypothesis as to why it is so complicated, and describe a simplified protocol with very similar features.

* Internal Accession Date Only

Personal use of this material is permitted. Permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or distribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE. ¹Hamburg University of Technology

[♦] IEEE Computer Security Foundations Symposium, 6-8 July 2007, Venice, Italy

© Copyright 2007 IEEE

Approved for External Publication

Electing the Doge of Venice: Analysis of a 13th Century Protocol*

Miranda Mowbray
HP Laboratories, Bristol
miranda.mowbray@hp.com

Dieter Gollmann
Hamburg University of Technology
diego@tu-harburg.de

Abstract

This paper discusses the protocol used for electing the Doge of Venice between 1268 and the end of the Republic in 1797. We will show that it has some useful properties that in addition to being interesting in themselves, also suggest that its fundamental design principle is worth investigating for application to leader election protocols in computer science. For example, it gives some opportunities to minorities while ensuring that more popular candidates are more likely to win, and offers some resistance to corruption of voters.

The most obvious feature of this protocol is that it is complicated and would have taken a long time to carry out. We will also advance a hypothesis as to why it is so complicated, and describe a simplified protocol with very similar properties.

1. Introduction: the protocol

The 1268 protocol for the election of the Doge of Venice was used with only minor changes for over five centuries, until the fall of the Venetian Republic in 1797. Descriptions of the protocol appear in books by Tappan ([19] pp.51-54) and Norwich ([15] pp.166-167).

The protocol was in ten rounds, the first nine of which produced an electoral college for the next round. The college for the first round was the entire electorate—the members of the Great Council of oligarchs aged 30 or over. No two members of the same family were allowed in the same college. Each round was one of two different types. In the

first type of round, the college for the next round was drawn by lot from the current electoral college. In the second type of round, the current college elected the next college, and every oligarch in the next college had to be approved by a certain minimum number of members of the current college. The college sizes and minimum number of approvals required in each round are given in Table 1: n denotes the size of the entire electorate.

Table 1. College sizes and minimum approval numbers

Round	type	size of college	approvals
1	lot	n	—
2	lot	30	—
3	election	9	7
4	lot	40	—
5	election	12	9
6	lot	25	—
7	election	9	7
8	lot	45	—
9	election	11	9
10	election	41	25

Norwich [15] remarks that this protocol “strikes the modern mind as ridiculous”. However, in this paper we will make a case that it has some properties that are not ridiculous, and that make its general principle—that of repeatedly reducing an electoral college by lot and then increasing it by election—worth investigating for application in leader election protocols in computer science. We will show that the protocol offers opportunities to minorities while ensuring that more popular candidates are more likely to win; that it may offer some resistance to corruption; and that it appears to assist the emergence of compromise candidates (where this is necessary) by amplifying small advantages. These properties may have contributed to the extraordinary stability and longevity of the Venetian Republic.

The rest of the paper is structured as follows. Section 2 briefly introduces leader election protocols. The main results of the paper are in Section 3, which analyses the Doge

*Copyright (c) 2007 IEEE. Reprinted from 20th IEEE Computer Security Foundations Symposium (Venice, Italy, July 6-8 2007), pp. 295–308, ISBN 0-7695-2819-8/07. This material is posted here with the permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of HP Labs Bristol’s products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

election protocol, under some simplifying assumptions on voter preferences. We investigate the probability that a minority Doge is elected, the probability that different factions are forced to negotiate, the optimum faction size for a voter to belong to, and the resilience of the protocol, comparing the performance of the protocol with that of some other leader election protocols. We also report experimental results on the emergence of compromise candidates.

The protocol is randomized, in the sense that if an election is held twice with the same candidates and same voter preferences, two different candidates may win (but it is not purely random—popular candidates have a better chance of winning than unpopular ones.) Standard criteria for judging election methods are in general not suitable for randomized protocols. Section 4 generalizes several criteria to make them more suitable for judging randomized protocols, and applies these to the protocol.

The subsequent sections take a closer look at the historical specifics of the protocol’s operation. Section 5 investigates the issue of incumbent bias, using historical data on the protocol’s performance. Section 6 is concerned with the choice of college sizes and minimal approval numbers in the protocol. Section 7 describes an attack on the protocol in the way that it was originally carried out in Venice, which is however straightforward to protect against with today’s technology.

Finally, Section 8 discusses the most obvious feature of the protocol: that it is very complicated. We argue that having a complicated and lengthy protocol for the election of the Doge might have been beneficial as “security theatre”. We describe a simplified protocol with very similar properties, for applications in which complexity is not valuable.

2. Leader election protocols

Several computer science protocols require a step in which one computer, or one process, is elected by the system as the leader. For example, this may be used to maintain consensus in asynchronous systems [9] or for recovery after a network partition [6, 13]. In general the electing computers or processes are not mutually trustworthy. Although several more complex election procedures have been suggested in order to ensure efficiency or deal with crashes (see e.g. [1]), less attention has been paid to the problem of untrustworthiness, and in practice the election is often done by simple majority voting. The leader plays an important part in the protocol, and may have more security privileges or denial-of-service opportunities than non-leaders. A malicious computer owner might therefore attempt to subvert or corrupt other electing computers in order to make a system under his control the leader in the protocol.

Simple majority voting has the drawback that an attacker who corrupts just over half the voters with a virus attack

or because of a common security loophole can be sure to gain leadership. As a result of the degree of similarity of modern computer systems and their configurations, it is not unreasonable to assume that some security attacks will be successful on a majority, but not all, of the voters. Clearly it is sensible that candidates supported by many voters should be more likely to become leader than candidates supported by few voters. However, if there is a nonzero probability that a minority candidate is elected, then such an attack will not succeed every time.

This points to a general resilience property for voting protocols: if an attacker wishes to have a probability at least c of obtaining a required result, what proportion of the voters does the attacker need to corrupt to ensure this? This resilience property is one of several properties considered in the next section.

3. Analysis

This section investigates some properties of the protocol for the election of the Doge. The protocol was certainly not designed with exact knowledge of these probabilistic properties, because probability was not well understood in the 13th century. However, one can make a reasonable case that the oligarchs of Venice may have intended similar properties to hold, and evolved the protocol by trial and error until they did. Simpler versions of the protocol were used to elect the Doge before 1268, and the day-to-day business of Venice was carried out by committees of various sizes drawn by lot or elected from electoral colleges of various sizes, so the oligarchs had considerable experience of the results of these operations.

3.1. Related work

As far as we are aware, just two other papers investigate the mathematical properties of the protocol. Lines [11] discusses approval voting, which is the method of voting used in each round which increased the college size, and in the final round. Candidates for the next college (or for the Dogeship, in the case of the final round) were proposed, and a ballot was held in which the current college members signalled either their approval or disapproval of each candidate, with no limit on how many candidates they approved or disapproved. Candidates receiving the required minimum number of approval votes joined the next college. If not enough candidates received the minimum, the college repeated the process, holding another ballot.

Lines points out that provided the votes are performed concurrently for each candidate, this system has the advantage over single-vote systems that there is no need for tactical voting: supporters of minority candidates can approve them without increasing the likelihood that candidates that

they dislike will be elected. If on the other hand each candidate is considered by the college in turn, Lines points out that tactical voting may be necessary. In 1268 the protocol specified that each candidate was considered in turn, but concurrent voting was introduced at a later date—it is not clear exactly when.

In fact, there is a case (described in [15] p.300) in which tactical voting allegedly determined the outcome of one election for Doge. In 1423 Francesco Foscari, an underdog candidate, received 17 approval votes out of 41 in the ninth ballot by the final college and 26 approval votes in the tenth ballot, thus winning the election. It was claimed that his supporters had engineered this win by voting in earlier ballots for a candidate that no-one wanted, thus enticing others to vote for Foscari, and then suddenly switching their votes. Presumably in 1423 concurrent voting had not yet been introduced.

Coggins and Perali [4] look at the minimum approval numbers used in the protocol. They point out the remarkable fact that 25 is *exactly* the minimum approval number that should be chosen for the final round in order to make the protocol satisfy Cablin and Nalebuff's 64% majority rule [3]. This has the effect that, under some plausible assumptions on the way that the voters form their preferences, there will not be any other oligarch who could have gained this number of approvals if he had stood against the selected Doge in a two-candidate election by the final college, and that the selected Doge is the only oligarch satisfying this property.

Neither Lines nor Coggins and Perali have an explanation why the protocol should have more than two rounds—one using lot-drawing, to make it impossible for those interested in influencing the result (by bribing electors, for example) to forecast the membership of the final college, and one using election, to take into account the views of the electorate.

3.2. Assumptions for the analysis

One of the assumptions that we will make for the analysis is that there would be few candidates in an election for Doge. Theoretically the number of candidates was equal to the number of voters—any oligarch could become Doge. In practice, however, voters clustered in factions supporting a particular oligarch for Doge. Indeed, the 75 Doges elected in the five centuries in which this protocol was used have only 44 surnames between them, demonstrating the dominance of a relatively small number of powerful families.

The number of families of oligarchs in Venice in the late Thirteenth Century was 206. In 1297 the list of these families was “closed”, making it rarer for new families to obtain the right to sit on the Great Council. The total number of families in the Great Council between that date and the end

of the Venetian Republic was 532. However, by comparing the list of Doges supplied by Norwich ([15], pp.641-2) with the data on the membership of the Great Council supplied by Raines ([16], Appendix 1) it can be seen that 70 of the 75 Doges elected using this protocol came from families who had a seat on the Great Council in 1297. This is particularly striking given that according to Raines, 22.8% of these families had died out by 1400. Moreover, over a third of the 75 Doges came from the select group of twenty-four families which were traditionally considered to be the founding families of Venice. (In fact, they came from just thirteen of these families.)

Apart from the dominance of a few families, another reason why there might be few candidates in practice is that the Dogeship was not necessarily a very enviable position. The Doge had to pay state expenses from his own pocket. There were laws strictly limiting his actions and ensuring that he could not profit financially from his role.

This has a parallel in leader election protocols for computers; the computer that wins the election will tend to have to contribute more resources to the operation of the protocol for which it is leader than non-leader computers do.

Interestingly, if an oligarch thought that there was a danger that he might be elected Doge against his wishes, there is a strategy that he could follow to try to ensure that he did *not* become Doge which is quite similar to the strategy to become Doge. In both cases, he should collect as large a faction as he could. In all colleges but the final one, his faction would vote for other members of the same faction. The only difference is in the final college, where his faction would vote for him if he wished to be Doge, and for someone else if he did not wish to be Doge. However, in this paper we will assume that if an oligarch has a (non-empty) faction supporting him, then he does wish to become Doge.

In order to avoid speculating on the results of negotiations between different factions, we will initially limit the analysis to the case where the electorate can be divided into just two factions, each supporting a different candidate; and that voters do not move between one faction and the other during the election process. We will further assume that if the number of members of one of the factions in a particular electoral college is more than the minimum approval number for the election by this college, then the members of this faction will elect as many members of the same faction as possible to the next college; and if on the other hand the number of members of each of the factions in the college is smaller than the minimum approval number, then the proportion of the members of the different factions in the next college is chosen to be as close as possible to the proportion in the current college. (We need to make some assumption about what happens in this last case, because if neither faction has the minimum number of approvals required then the two factions have to negotiate to decide the membership

of the next college. In Subsection 3.5 we will explore how likely it is that this situation will occur.)

To be precise, suppose the faction with most members in the current electoral college has f_1 such members, and f members in total, and that the sizes of the current college and the next college are c, c' , with $c' > c$. Write f_2 for the number of members of this faction that are elected to the next college. If f_1 is greater than or equal to the minimum approvals number for the current college, then we assume that $f_2 = \min\{f, c'\}$, and if f_1 is less than the minimum approvals number then we assume that $f_2 = \min\{f, \lfloor f_1 \cdot c'/c \rfloor\}$ with probability $1 - \text{frac}(f_1 \cdot c'/c)$, $f_2 = \min\{f, \lceil f_1 \cdot c'/c \rceil\}$ with probability $\text{frac}(f_1 \cdot c'/c)$, where frac denotes the fractional part.

We assume that the final college chooses a Doge who has the approval of a majority of the college.

Finally, in the absence of information on the distribution of family sizes, we will ignore effects arising from the rule that no two members of a college could be from the same family.

Clearly this protocol can be applied for any number of voters over 44. We will investigate its properties when the number of voters is n , which is assumed to be over 44. When we have to choose some value for n in order to plot a figure, we will choose the value 480, which was the number of Venetian oligarchs in 1268.

3.3. Effect of a single round

Under the assumptions introduced in the previous subsection, for any round of the protocol before the final round it is possible to calculate precisely the probability distribution of the number of members of a faction in the next college, given the probability distribution for the college for this round.

Let c, c' be the sizes of the current college and the next college. Let n be the size of the total electorate and f be the total number of members in the faction, and for $0 \leq i \leq f$ let p_i, p'_i be the probabilities that the current college and the next college respectively contain exactly i members of the faction.

First, suppose that the round is of lot type. Then for $0 \leq i \leq f$, p'_i is equal to

$$\binom{c}{c'}^{-1} \cdot \sum_{j=i}^f \binom{j}{i} \binom{c-j}{c'-i} p_j$$

Now suppose the round is of election type. Write $N(i)$ (which depends on n, m, f, c and c' as well as i) as shorthand for the set of integers j satisfying

$$n - m < j < \min\{m, f\}, \lfloor c'j/c \rfloor = i$$

Then p'_i is equal to

$$\sum_{j=0}^{\min\{n-m, f\}} p_j + \sum_{j=n-m+1}^{\min\{m, f\}-1} \text{frac}(c'j/c) p_j$$

if $i = 0$,

$$\sum_{j \in N(i)} \text{frac}(c'j/c) p_j + \sum_{j \in N(i-1)} (1 - \text{frac}(c'j/c)) p_j$$

if $0 < i < \min\{f, c'\}$,

$$\sum_{j=m}^f p_j + \sum_{j \in N(k), k \geq i} p_j + \sum_{j \in N(i-1)} (1 - \text{frac}(c'j/c)) p_j$$

if $i = \min\{f, c'\}$, and $p'_i = 0$ if $c' < i \leq f$.

The effect of the final round can also be calculated exactly. Suppose that for all $(0 \leq i \leq f)$, p_i denotes the probability that the faction has exactly i members in the final college, and c is the size of the final college. Then, under our assumption about the behaviour of the final college, the Doge elected is from this faction with probability

$$\sum_{j > c/2} p_j$$

if c is odd, and probability

$$\sum_{j > c/2} p_j + \frac{1}{2} p_{c/2}$$

if c is even.

3.4. Probability of a minority Doge

If a faction has f members, this means that the college for the first round of the protocol (which is the entire electorate) contains f members of this faction with probability 1. By repeatedly using the formulae in Subsection 3.3, it is therefore possible to calculate the probability that a faction member is elected Doge. The formulae can also be used to calculate the probability that a faction member is elected Doge if the election protocol is altered to become a truncated protocol in which the final round of voting for the Doge is done by the college for a round earlier than the tenth. Figure 1 shows the probability that the candidate supported by less than half the electorate becomes Doge, for the full protocol and various truncated protocols in an electorate of size 480.

The horizontal axis shows the fraction of the electorate that support the minority candidate, and the vertical axis the probability that this candidate is elected. The different lines show the results if the final round of voting for the Doge is




Figure 1. Probability of electing a minority Doge, for the protocol truncated at rounds 2,4,6,8, and not truncated: lower lines correspond to earlier truncation

done by the electoral college for the second, fourth, sixth, eighth or tenth round. The 1268 protocol has ten rounds, and the other three lines show the results of truncated protocols. The line for each protocol lies above those for the protocols which are truncated versions of it, so the effect of additional rounds of voting is to increase the chances of minority candidates becoming Doge. However, as can be seen, the effect of each additional pair of rounds is smaller, and a 12-round protocol is unlikely to have a very different effect from a 10-round protocol; so the oligarchs were sensible to stop the protocol at the 10th round.

This result shows that the protocol offers some support to minorities—in contrast to a simple majority protocol, this protocol does offer the possibility that a minority candidate becomes Doge—while ensuring that the most popular candidate is most likely to win.

Since we assume that the Doge is supported by a majority of the final college, which has 41 voters, a candidate with fewer than 21 supporters cannot be elected. In some modern voting systems using proportional representation to afford some protection of minorities (such as the electoral systems in Germany, Russia and New Zealand), representation is only offered to parties with at least 5% of the votes [17]. Five percent of an electorate of 480 voters is 24 voters.

3.5. Probability of the election forcing a negotiation between factions

It is possible that during the election process there is some college in which neither faction has as many members in the college as the required minimum number of approval votes for the election by that college. When that happened,

the members of the college in the two factions would have had to negotiate with each other in order to decide whom to elect to the next college. If either faction was not satisfied, they would have the power to stall the election indefinitely. (In practice, no election of a Doge stalled.) This feature offers a protection to minorities in addition to the possibility of the election of a minority Doge; sizeable minorities would be quite likely to reach a position during the election process during which they had some power to negotiate favourable treatment from the majority, even if the majority candidate won.

The probability of this happening can be calculated by repeatedly using the results of Subsection 3.3. The result for $n = 480$ is plotted in Figure 2.




Figure 2. Probability that at some stage during the election process there is a college in which the factions have to negotiate

Note that for $n = 480$ it is more likely than not that there is such a college during the election provided that the minority candidate has at least 138 supporters, about 29% of the electorate.

If the size f of the minority faction is at least 43, then there will be such a college at some point during the election if and only if the college for the third round is of this type. It follows that if $f \geq 43$ there is a simple formula for the probability that the election process forces a negotiation between the factions: it is equal to

$$\sum_{i=3}^6 \binom{f}{i} \binom{n-f}{9-i} \binom{n}{9}^{-1}$$

Our assumptions about the behaviour of voters in such a college can be dropped without affecting this formula. If there are more than two factions, the probability of this occurring during the election is of course even higher.

For the potential application to leader election protocols in computer science, this feature is not very useful, because

software processes are not (yet) as skilled at negotiation as Venetian oligarchs were. Our suggestion is that for this application, when there is a college in which no faction contains as many members as the minimum number of votes needed for election by the college, the membership of the next college should be chosen as close to proportionally as possible; this does not require any negotiation.

3.6. Expected gain by voters

There were strict laws limiting the actions of a Doge and making it difficult for him to use his office to enrich himself. However, it is plausible that a successful candidate for Doge would reward the members of his faction. Suppose the Doge, if elected, divides a fixed bounty between the other members of his faction. An interesting metric to calculate is the expected fraction of the bounty gained by a voter in a faction of a particular size whose candidate would do this if elected. This is the probability that a faction of that size elects their candidate, divided by one less than the faction size. (One less, because the candidate is part of the faction, but does not receive any of the bounty.) The larger the value of this metric for a particular faction size, the more attractive it is to be a member of a faction of this size.

Figure 3 shows the results for three election protocols; simple majority (with ties resolved randomly), probabilistic proportional representation (in which all voters place a slip of paper naming their preferred candidate in an urn, one is selected at random, and the candidate that it names wins the election), and the protocol actually used for the election of the Doge. In the rest of this paper we will denote these protocols by $\mathcal{P}_M$, $\mathcal{P}_P$ and $\mathcal{P}_D$ respectively.




Figure 3. Expected gain by a voter in a faction of a given size

For $\mathcal{P}_P$, it turns out that the more unpopular a candidate is, the greater the expected gain by their supporters. For

$\mathcal{P}_M$, the voters with greatest expected gain are those who support a candidate with only just over half the votes. So under either of these two protocols there is a possible danger that supporters of a popular candidate might try to dissuade other voters from supporting him. The protocol $\mathcal{P}_D$ has the property that the expected gain is largest for supporters of candidates whose factions contain 357 oligarchs, which is about 3/4 of the electorate. Thus, this protocol encourages the building of broad factions. As the number of rounds in the election increases, the faction size giving the highest expected gain to faction members increases.

3.7. Resilience

In the election for the Doge, voters were required to act for the general social good (or, at any rate, for the good of the oligarchs of Venice). Before the final college voted, each member of the college had to each swear an oath that he would act for the good of the Republic. However, it is plausible that a voter might be corrupted into voting in favour of a candidate that he thought might not be the best for the Republic in general. In return for a bribe or personal favour, he would join the candidate's faction: that is, in rounds other than the final one he would approve an oligarch for the next college if and only if the oligarch was a member of the candidate's faction, and in the final round he would vote for the candidate.

In this subsection we give a formal way of comparing the resilience of different protocols, and use it to compare the resilience of the three election protocols investigated in the previous subsection. The question considered is this: suppose that an attacker wishes to have probability at least d that his favoured candidate becomes Doge. What percentage of the electorate does he need to corrupt to obtain this? Intuitively speaking, the higher this percentage is, the more resilient the protocol is.

In general, the probability that the attacker's candidate will win depends on what proportion of the electorate will support the attacker's candidate without being corrupted, and the strategies that the uncorrupted voters follow, as well as on d and on the percentage of voters that are corrupted. For some protocols it will depend also on the exact choice of which voters are corrupted, rather than just the percentage. It will also depend on the size of the electorate, although for sensible protocols the answer will not vary much with the size of the electorate provided that the electorate is large. In order to deal with these variations, we measure the maximum percentage of the voters that the attacker needs to corrupt over all possible strategies by uncorrupted voters, in the limit as the electorate tends to infinity, assuming that the attacker corrupts the best possible set of voters to fulfill his goal.

Several researchers have investigated the resilience of

leader election protocols. Antonakopoulos [2] gives the following typical definition of resilience (among several equivalent variants):

Definition 3.1 (ϵ -resilience) Let $\epsilon > 0$ be a constant independent of the number n of voters. A leader-election protocol $\mathcal{P}$ is called ϵ -resilient for $t = b(n)$ if and only if for all sufficiently large n ,

$$\text{fail}_{\mathcal{P}}(n, b(n)) \leq 1 - \epsilon$$

where $\text{fail}_{\mathcal{P}}(n, t)$ is the probability that an attacker who corrupts the best possible set of t voters to fulfill his goal will fulfill it.

Antonakopoulos and other researchers in this area focus on distributed elections where the uncorrupted voters behave randomly, and where attackers and corrupted voters may try to influence the outcome of the election by tuning their “random” inputs. In the Venetian election, uncorrupted voters in the same faction cooperate with each other, and the lot-drawing during the protocol is the source of the randomness. So results previously proved about the resilience of leader election protocols under the assumption that the participants are the source of randomness do not necessarily apply for the election protocols we are studying here. However, we can still use the definition above to compare the resilience of $\mathcal{P}_M$, $\mathcal{P}_P$, and $\mathcal{P}_D$.

In order to achieve a probability at least d of having his favoured candidate elected Doge under the protocol $\mathcal{P}$, the number of candidates that an attacker will need to corrupt (for any choice of strategies by the uncorrupted voters, in a large electorate) is $q(\mathcal{P}, d) \cdot n$, where

$$q(\mathcal{P}, d) = \sup\{q : \mathcal{P} \text{ is } (1 - d)\text{-resilient for } t = q \cdot n\}$$

If $\lim_{n \rightarrow \infty} \text{fail}_{\mathcal{P}}(n, q \cdot n)$ is a strictly increasing function of q —which is true if $\mathcal{P}$ is one of the protocols $\mathcal{P}_P$, $\mathcal{P}_D$ —then $q(\mathcal{P}, d)$ is just equal to the q satisfying

$$\lim_{n \rightarrow \infty} \text{fail}_{\mathcal{P}}(n, q \cdot n) = 1 - d$$

If D is an interval of $[0, 1]$ and $\mathcal{P}_1$, $\mathcal{P}_2$ are election protocols, we say that $\mathcal{P}_1$ is *more resilient than* $\mathcal{P}_2$ on D if and only if $q(\mathcal{P}_1, d) > q(\mathcal{P}_2, d)$ for all $d \in D$.

For all protocols $\mathcal{P}$, $q(\mathcal{P}, 0) = 0$. In the case that $q(\mathcal{P}, d)$ is a continuous function of $d \in [0, 1]$ and that the probability of an attacker’s candidate winning does not depend on the precise identity of the corrupted voters, but only on their number, we have that $q(\mathcal{P}, d) + q(\mathcal{P}, 1 - d) = 1$. It follows that if two protocols $\mathcal{P}_1$ and $\mathcal{P}_2$ both satisfy this case, and $\mathcal{P}_1$ is more resilient than $\mathcal{P}_2$ on an interval $[c_1, c_2]$ with $0 < c_1 \leq c_2 < 0.5$, then $\mathcal{P}_2$ is more resilient than $\mathcal{P}_1$ on $[1 - c_2, 1 - c_1]$. The same result holds if $\mathcal{P}_2$ is replaced by simple majority voting (with a coin-toss for a tie). When

comparing different protocols, therefore, it is not reasonable to expect one protocol to be more resilient than the other on the entire interval $[0, 1]$. If the protocols satisfy the conditions described above, in general one will be more resilient over some proper sub-interval of $[0, 0.5]$, and the other more resilient in the mirror sub-interval of $[0.5, 1]$. The choice of protocol will depend on whether it is more important to prevent attackers from gaining small advantages easily, or to hinder attackers who have the ability to corrupt a large proportion of the electorate.

We will now give results for the three election protocols $\mathcal{P}_M$, $\mathcal{P}_P$, $\mathcal{P}_D$ considered in the previous section. It is straightforward to check that $q(\mathcal{P}_M, d)$ is zero for $d = 0$ and 0.5 otherwise, and $q(\mathcal{P}_P, d) = d$ for all d in $[0, 1]$. The value of the function $q(\mathcal{P}_D, d)$ does not have such a simple expression: it is the value q satisfying

$$\sum_{i=0}^9 [a_i \cdot \binom{9}{i} q^i (1 - q)^{9-i}] = d$$

where a_i is the probability that a faction with f members and exactly i members in the college for round 3 will have one of their members elected Doge, if $\min\{f, n - f\} \geq 43$. (This probability is the same for all values of f and n satisfying $\min\{f, n - f\} \geq 43$). The values of a_i can be calculated using the results of Subsection 3.3: they are $a_0 = a_1 = a_2 = 0$, $a_3 \approx 0.1955$, $a_4 \approx 0.3929$, $a_5 \approx 0.6071$, $a_6 \approx 0.8045$, $a_7 = a_8 = a_9 = 1$.

The functions $q(\mathcal{P}_M, d)$, $q(\mathcal{P}_P, d)$ and $q(\mathcal{P}_D, d)$ are plotted in Figure 4. On the interval $[0.0001, 0.4999]$, $\mathcal{P}_D$ is more resilient than $\mathcal{P}_P$ and less resilient than $\mathcal{P}_M$. On the mirror interval $[0.5001, 0.9999]$, $\mathcal{P}_D$ is less resilient than $\mathcal{P}_P$ and more resilient than $\mathcal{P}_M$.

3.8. Finding a compromise

If one does not assume that the oligarchs were divided into two main factions only, but that there was a set of eligible candidates supported by groups of various sizes, then the election process would somehow have to arrive at a compromise candidate. We assume in addition that there is sufficient familiarity between the oligarchs so that their preferences were known in principle. To analyse how a compromise might emerge in this election process, we have modelled the election process as follows.

Each oligarch has an ordered list of ten preferred candidates. Let c_{ij} be the candidate in position j in the list of oligarch i . When a new college is elected, the members of the electing college pick those oligarchs whose preferences are closest to their own. Assuming that they hold stronger views about candidates at the top of their list, we attach the weight $\frac{1}{2^j}$ to position j . The similarity $\sigma(i, m)$ between the




Figure 4. Fraction of voters in a large electorate that it is necessary for an attacker to corrupt to ensure a given probability of getting his candidate elected.

preferences of oligarchs i and m is given as

$$\sigma(i, m) = \sum_{j=1}^{10} \sum_{k=1}^{10} \delta(c_{ij}, c_{mk}) \frac{1}{2^{j+k}}.$$

where $\delta(c_{ij}, c_{mk}) = 1$ if $c_{ij} = c_{mk}$ and zero otherwise. To decide the members of the next college we add the similarity scores of all electors and take the required number of members with the highest overall score. (In this computation, similarities $\sigma(i, i)$ are not included.) The last college C elects the candidate X who maximizes

$$\sum_{i \in C} \sum_{j=1}^{10} \delta(c_{ij}, X) \frac{1}{2^j}$$

Using this model, we performed the following experiment. We assume that all oligarchs are eligible. All positions in the preference vectors are selected at random (equal distribution) but for any fixed i the entries c_{ij} have to be different. Differences in the popularity of candidates are thus only the consequence of this random selection. We then run the election 10000 times and record which candidate was elected how often. The 10000 runs are all made with the same preference vectors: the source of variation in the winning candidate is just the drawing of lots that takes place during the election. We have run this experiment with three different sets of preference vectors and found that in each case there was one candidate who was chosen in 2000–2500 of the election runs (see fig. 6 for the results of one such experiment). The initial advantages of the candidates, com-




Figure 5. Initial weighted preferences per candidate.

puted for each candidate X as

$$\sum_{i \neq X} \sum_{j=1}^{10} \delta(c_{ij}, X) \frac{1}{2^j}$$

are given in fig. 5. This might suggest that the election process had the property of amplifying small advantages in an initially more chaotic preference landscape.

We have also examined the influence of the number of rounds on the result. We have run our experiment for a six round protocol where rounds 5–8 have been omitted, an eight round protocol where rounds 7–8 have been omitted, and a twelve round protocol where rounds 7–8 are repeated. The election results given in figure 7 show that the move from a six round protocol to an eight round protocol significantly increases the advantage of the favourites, but the move to a ten round or twelve round protocol has little impact on the expected outcome. Hence, also this experiment suggests that there were merits in not going beyond ten rounds.

4. Criteria for randomized election protocols

In this section we will suggest some general criteria for judging randomized election protocols, and apply them to the Doge election protocol. Since standard criteria for judging election protocols do not assume that there are only two factions, in this section we will drop that assumption, and consider the behaviour of the Doge election protocol when there may be more than two factions. We will however retain our assumptions from Subsection 3.2 on how faction members vote.

Cretney’s resource site for election methods [5] lists twelve criteria for judging election protocols; these criteria are given in the Appendix to this paper. Not one of the




Figure 6. Frequency of wins per candidate in the actual ten round protocol.




Figure 7. Frequency of wins per candidate in 6, 8, 10, 12 round protocols, logarithmic scale.

twelve criteria holds for the Doge election protocol. The problem is that the criteria are not suitable for judging randomized protocols. The criteria generally take the form “If X wins, and some particular changes are made to the voter preferences or set of candidates, then X still wins”; however for a randomized protocol, such as the Doge election protocol, if X wins and *no* changes are made to the voter preferences or set of candidates, then X may not win next time.

Cretney’s resource site also lists election protocols, and two of those listed are randomized (“Random”, in which every candidate has an equal probability of winning, and “Random ballot”, which is $\mathcal{P}_P$). Randomized election protocols have the drawback that it may be more difficult to detect when they have been incorrectly implemented. However their use for important purposes is certainly not confined to 13th century Venice—for example, jury selection systems are usually randomized. There appears therefore to be a need for criteria for judging randomized election protocols.

It is possible to generalize each of the criteria listed by Cretney so that they can be sensibly applied to randomized protocols. In most of the cases it is enough just to replace “ X wins” with “ X is the candidate with the highest probability of winning” and “ X loses” by “ X is not the candidate with the highest probability of winning”. This gives straightforward generalizations of the Majority, Consistency, Pareto, Secret Preferences, Concordet, Concordet Loser, Independence of Clones, and Reversal Symmetry criteria. Since they are straightforward we will not give the definitions of these generalised criteria here.

In order to generalise the Smith criteria and Local Independence from Irrelevant Alternatives, it is useful to introduce the idea of a *generalised Smith set*. This is the smallest nonempty set of candidates such that for all Y in the set and Z not in the set, Y has a higher probability than Z of winning the election if all candidates but Y and Z are eliminated. Note that the set of all sets with this property is nonempty and nested; it follows that the generalised Smith set always exists, and is unique. It is equal to the usual Smith set if the election protocol is not randomized.

The Smith, Local Independence from Irrelevant Alternatives, Monotonicity and Mutual Majority criteria can be generalised to the following:

Criterion 4.1. (Generalised Smith) *If X has the highest probability of winning, X must be a member of the generalised Smith set.*

Criterion 4.2. (Generalised Local Independence from Irrelevant Alternatives) *Suppose X is the candidate with the highest probability of winning, a new candidate Y is added, and Y is not in the generalised Smith set. Then X*

is still the candidate with the highest probability of winning.

Criterion 4.3. (Generalised Monotonicity) *If X is not the candidate with the highest probability of winning, and some of the voters change their mind so as to rank X lower than they did before, then X is still not the candidate with the highest probability of winning.*

Criterion 4.4. (Generalised Mutual Majority) *If there is a set of candidates for which a majority of voters rank any candidate in the set higher than any candidate not in the set, then it is more likely that a candidate in the set will win than that a candidate not in the set will win.*

Each of the generalized criteria reduces to the original case if the protocol is not randomized. They can therefore be used to compare protocols that are randomized with ones that are not in a sensible fashion, as well as to compare different randomized protocols.

We will now determine which of the generalized criteria are satisfied by $\mathcal{P}_M$, $\mathcal{P}_P$ and $\mathcal{P}_D$.

Observe that for the protocols $\mathcal{P}_M$ and $\mathcal{P}_P$, and also for $\mathcal{P}_D$ under the assumptions that we have made on how faction members vote, candidate X has a higher probability of winning than any other candidate if and only if for every candidate Y not equal to X , the number of voters that rank X higher than any other candidate is greater than the number of voters that rank Y higher than any other candidate. We will refer to this as the *highest rank property*.

It follows directly from the highest rank property that the three protocols satisfy the generalised versions of the Majority, Consistency, Pareto, and Secret Preferences criteria. For the generalised version of Monotonicity, observe that when the voters change their mind the number of voters that rank X highest cannot increase, and the number of voters that rank other candidates highest cannot decrease; so again all three protocols satisfy this because they satisfy the highest rank property.

Next, consider the following example. There are three candidates X_1, X_2, X_3 with faction sizes 200, 180, 100. The voters who are not in X_1 's faction consider X_1 their least favourite candidate, so would join the remaining faction if their candidate were eliminated. The voters in X_1 's faction consider X_2 their least favourite, so would join X_3 's faction if X_1 was eliminated.

Under any protocol satisfying the highest rank property, the candidate with the highest probability of winning the election for this example is X_1 if no candidate is eliminated; X_3 if X_1 or X_2 is eliminated; and X_2 if X_3 is eliminated. The set $\{X_2, X_3\}$ is a clone set, and the generalised Smith set is $\{X_3\}$. It can be checked that with appropriate choices for X , this furnishes a counterexample to the generalised versions of Concordet, Concordet loser, Inde-

pendence of Clones, Reversal Symmetry, Smith, and Local Independence from Irrelevant Alternatives for any protocol satisfying the highest rank property.

This leaves the Generalised Mutual Majority criterion. It is easy to see that this property is satisfied by $\mathcal{P}_P$. However, it does not hold for all protocols satisfying the highest rank property. For a start, it does not hold for $\mathcal{P}_M$. As an extreme example, suppose that there are 242 candidates $X, X_1, \dots, X_{241}$ such that all of $X_1, \dots, X_{241}$ have a faction size 1 (so these candidates are supported by no-one but themselves) and regard X as their least favourite candidate, whereas the faction for X consists of all 239 other voters. If Generalised Mutual Majority held, then X should have a probability less than 0.5 of winning—but under $\mathcal{P}_M$, X wins with probability 1.

In fact, X also wins with probability more than 0.5 under $\mathcal{P}_D$, so $\mathcal{P}_D$ does not satisfy Generalised Mutual Majority either. The reason that X has a probability greater than 0.5 of becoming Doge under $\mathcal{P}_D$ is that the situation can arise that the majority set $X_1, \dots, X_{241}$ have enough members in a college to give the minimum required number of approval votes (and thus ensure that one of their candidates became Doge) if they all voted for the same candidate, but they do not, because they all have different candidates as favourite, and the election is made proportionally to the faction sizes in the electing college, thus leaving X with a chance of being elected. The faction supporting X does not have this problem—since it is only a single faction, if it has enough members in some college formed during the election process to give the required number of approval votes, it will make sure that X is elected. This difference in powers of coordination between X 's faction and the majority set is enough to outweigh the fact that the majority set has two more members than X 's faction.

5. Incumbent bias

A possible reason for designing an election protocol to have multiple rounds is to avoid favouring an incumbent. In elections where the incumbent is allowed to stand, voters may have some psychological inhibitions about voting against an incumbent in a single-round election, but multiple-round elections may allow strong alternative candidates to emerge in discussions.

Doges were elected for life (although a few were made to resign before they died). So in an election for Doge, the incumbent could not stand. Election for life could be viewed as an extreme case of favouring the incumbent. On the other hand, according to Maranini ([12], p.274), older candidates were preferred by voters in order to mitigate the risk of elections for life. Indeed, the Doges elected under this protocol did not tend to last long. The mean and median numbers of years served—measured as the date of death or resigna-




Figure 8. Lengths of service of Doges

tion minus the date elected—by these 75 Doges were 6.85 and 6 respectively. (This compares to mean and median reign lengths of 19.85 and 14 years for the English monarchs whose reigns began within the time period during which this protocol was used in Venice). In contrast, the mean and median lengths of service of the 43 Doges before 1268 were 11.35 and 11 years. Thus, if it is correct that the oligarchs wished to reduce the incumbent advantage by electing Doges who were unlikely to remain for many years, it appears that the 1268 protocol was more effective than the previous, simpler election protocols at reflecting this wish of the electorate.

Figure 8 shows the lengths of service of the Doges in order of election; the Doges from number 44 onwards were elected using the 1268 protocol. It can be seen that Doges with 5 or fewer years of service occur more frequently after the introduction of this protocol.

The Doge with the longest service, 34 years, was elected post 1268; interestingly this was Francesco Foscari, the alleged beneficiary of tactical voting mentioned earlier in this paper. It can be speculated that if the revised election rule using concurrent voting during ballots had been introduced before this election rather than after it, the election might have been won instead by the admiral Pietro Loredan, considered the most likely candidate for the Dogeship, who died 16 years after the election.

Another way of thinking about incumbent bias in the election of the Doge is to consider not the issue of the same oligarch winning two consecutive elections (which was impossible), but the issue of the election of an oligarch from the same family as the previous Doge. This happened twice in the 75 elections after 1268; in each case the two successive Doges were brothers. In contrast, seven of the 43 pre-1268 Doges came from the previous Doge's family, and—ominously for a Republic—in six of these cases the Dogeship passed from father to son. The reason for the reduction

of the advantage to the incumbent family appears to be the abolition in the mid-11th century of the practice of allowing Doges to appoint a co-regent, who was effectively his designated successor ([15], p.66). No Doges from the incumbent family were elected between then and 1268. However, the large number of rounds of the 1268 protocol may have helped to ensure that a strong incumbent bias did not creep back after this date.

The election of two pairs of successive Doges from the same family after 1268 might have been the result of residual incumbent bias, but it is also consistent with there being no bias in favour of (or against) the incumbent family, but a strong bias limiting the number of families from which the Doge might be elected. If each of the 75 Doges after 1268 had been independently selected in such a way that each of the 44 families which produced Doges had an equal chance of being the family of the Doge selected, and no other families had a chance, then it would have been (just) more likely than not that at least two of the Doges selected would be from the same family as the Doge preceding them. If the number of families with a chance had been 206—the number of oligarch families in the late 13th century—rather than 44, there would still have been a 5% chance of two or more such Doges being selected. (We are not suggesting that the Doges were selected this way. This calculation is only intended to show that the data do not imply an incumbent bias after 1268.)

6. Why these numbers?

Lines ([11], p.157) says about the college sizes and minimum approval numbers for the 1268 protocol that “After some effort to find a logical or probabilistic structure behind the numbers used in the scheme, I finally attributed them to the whims of Zorzi and/or his cohorts.” (The 1268 protocol was designed either by Ruggero Zorzi, or by Zorzi and others.)

In an earlier version of the protocol, used in 1178, the assembly of all the voters elected a college of 4, who elected a college of 40, who elected the Doge by majority voting. The minimum number of approval votes required for the election by the college of 4 was 3. In all the rounds but the final one of the 1268 protocol, the minimum number of approval votes required is $\lceil 3c/4 \rceil$, where c is the current college size. As far as we are aware, no-one else analysing this protocol has noted this simple formula, which is a straightforward generalization of the value of 3 for a college of 4. It has the general property that if the college is split into r factions of sizes $f_1 \geq f_2 \geq \dots f_r$, then no two or three factions can unite to obtain the required number of approval votes unless at least one of these factions has size f_1 ; so it is not possible for a pair or triplet of allied smaller factions to outvote the largest faction. In fact, for the college sizes used in the pro-

tocol, it is not possible even for an alliance of four smaller factions to outvote the largest faction.

The number of approval votes required for the election in the final round of the protocol, 25, is less than $3/4$ of the college size of 41, perhaps reflecting the difficulty of reaching a decision in a relatively large college. It is equal to $\lceil 3c/5 \rceil$ where c is the college size, although it may not have been selected for that reason. As mentioned in section 3.1, Coggins and Perali [4] have pointed out that under certain conditions on voter preferences the value 25 is exactly the one that makes the system as stable as possible.

We have experimented with some different sequences of college sizes. In order for the results after 10 rounds to be close to the results after 8 rounds, it is important for the colleges drawn by lot to be not too close in size to the colleges from which they are drawn. Apart from this effect there are broadly similar results for all the sequences we tried.

The college sizes may have symbolic meanings relating to religion or to the history of Venice; or, they may reflect sizes of known factions. If there was a powerful faction consisting of n families of oligarchs, the protocol might have been designed to include a college with size $2n + 1$ or with a requirement that the number of approval votes is at least $n + 1$. For example, the 25 approval votes required in the final round might have been chosen to ensure that in order to be elected a Doge would need the approval of at least one voter who was not a member of one of the 24 founding families of Venice.

The final college size of 40 from which the Doge was voted by majority in the 1178 protocol has an obvious drawback: the possibility of a tie. In 1229 the election for Doge did indeed reach a tie (which was finally resolved by a drawing of lots) and the election rules were amended to make the final college size 41. This final college size was retained for the 1268 protocol.

At first glance it appears odd that the protocol begins by picking 30 voters out of the entire electorate, and then 9 out of these 30, rather than directly picking 9 out of the electorate. We suspect that this was done to compensate for technical limitations of the lot-drawing process. The selection by lot was implemented by picking ballot balls from an urn. It would have been difficult to properly mix 480 ballot balls inside the urn, but much easier to do so for 30 ballot balls. Therefore the two-round process would have given a more random result than picking the 9 directly by drawing lots from an urn containing 480 balls.

7. An attack on the protocol

The easiest way to attack this protocol is probably through its random number generator, that is, the person who draws the lots. Originally the drawing was carried out by an oligarch appointed for the task. In 1328 this rule

was changed, presumably because it was noticed that there was a risk that the appointed oligarch might not draw fairly, and from then on the drawing of lots was carried out by a *balotino*, a boy who was selected as the first boy seen by the oligarch with a particular public position—the youngest member of the inner council of state—after this oligarch finished praying at St Mark’s Basilica on election day. (The title “balotino” is derived from “balota”, the Venetian word for the ballot ball used in the Doge’s election; the modern English word “ballot” is derived from the same Venetian word.)

Since the balotino was selected as the first boy to appear at a certain place and time known to everyone, it could be possible to train a boy to favour your faction when drawing lots, and to release this boy at just the right place and time to become the balotino.

To reduce this vulnerability, it would have been better to choose the place of selection of the balotino at the last minute, using a method relying on more than one oligarch for the choice. Moreover it might have been wise to introduce a process allowing oligarchs to raise an objection to a balotino who appeared to be biased.

For modern applications, there are services [14, 20, 7] which provide random numbers over the Internet, and in some cases also the code to produce them locally from a source of random noise. There are ways of combining “random” bit-strings from several different sources so that if any of the strings are random, then the resulting string will be random as well.

8. Why is the protocol so complicated?

The most obvious feature of the 1268 protocol is that it is complicated, and would have taken a long time to carry out. Francesco Da Mosto (cited in [10]) says that “Even we Venetians don’t understand the system”. Norwich [15] ends his description of the protocol with the comment that “With a system so tortuously involved as this, it may seem remarkable that anyone was ever elected at all.” (p.167).

Under the hypothesis that no voters change their mind during the election process, the process need not usually take as long as might appear. This is because if a faction supporting one particular candidate and containing at least 43 voters of different families gains at least as many members in some college as the minimum number of approval votes required for elections by this college, then it can be announced at once that this faction’s choice will be elected Doge, because the faction will be able to control the results of the elections by all subsequent colleges. This fact can be used to shorten the process when the voters are computers. However, it is plausible that in Venice no such shortening would have been allowed, so that the voters in the winning faction would have the opportunity to change their

minds between rounds, or to forge alliances advancing their second-choice candidates.

It has been suggested that the complexity of the protocol was an aesthetic choice by the Venetian oligarchs. Certainly the Venetian Republic produced some very complex and highly ornamented music and architecture. It is also possible that it was complex simply because no simpler protocol with the properties that were wanted had been found. We suspect however that this complexity served a particular function: that of security theatre.

8.1. Security theatre in the Venetian Republic

Schneier [18] has used the phrase “security theatre” to describe public actions which do not increase security, but which are designed to make the public think that the organization carrying out the actions is taking security seriously. (He describes some examples of this in response to the 9/11 suicide attacks.) This phrase is usually used pejoratively. However, security theatre has positive aspects too, provided that it is not used as a substitute for actions that would actually improve security.

In the context of the election of the Doge, the complexity of the protocol had the effect that all the oligarchs took part in a long, involved ritual in which they demonstrated individually and collectively to each other that they took seriously their responsibility to try to elect a Doge who would act for the good of Venice, and also that they would submit to the rule of the Doge after he was elected. This demonstration was particularly important given the disastrous consequences in other Mediaeval Italian city states of unsuitable rulers or civil strife between different aristocratic factions. It would have served, too, as commercial brand-building for Venice, reassuring the oligarchs’ customers and trading partners that the city was likely to remain stable and business-friendly.

After the election, the security theatre continued for several days of elaborate processions and parties. There is also some evidence of security theatre outside the election period. A 16th century engraving by Mateo Pagan depicting the lavish parade which took place in Venice each year on Palm Sunday shows the balotino in the parade, in a prominent position—next to the Grand Chancellor—and dressed in what appears to be a special costume (Kurtzman and Koldau [8], Figure 18).

8.2. A simplified protocol

In the context of leader election protocols in computer systems, it is more sensible to implement security theatre by, for example, publicizing proven results about the security of the systems and protocols used than by making

the protocols gratuitously complicated. We therefore offer a simplified version of the Doge election protocol.

The simplified protocol has seven rounds rather than ten, and its college sizes for rounds two to seven are 9, 45, 9, 45, 9, 45, with a minimum of 7 approvals out of 9 required in each round of election type: otherwise it is the same as the original. For all faction sizes the original and simplified protocols differ by less than 0.001 in the measures plotted in Figures 1 and 4, and by less than 0.0002 in the measures plotted in Figures 2 and 3.

9 Acknowledgements

Miranda Mowbray thanks Matthew Hennessey for informing her about this extraordinary protocol; Alessandra Ori for Venetian linguistic advice; the anonymous reviewers for their helpful comments; and HP Labs, for letting her study the habits of Venetian oligarchs as part of her job. Dieter Gollmann thanks Harald Sauff for his help in performing the experiments of Section 3.8.

References

- [1] M. K. Aguilera, C. Delporte-Gallet, H. Fauconnier, and S. Toueg. Communication-efficient leader election and consensus with limited link synchrony. *Proc. Principles of Distributed Computing (PODC’04)*, pages 328–337, July 2004.
- [2] S. Antonakopoulos. Fast leader-election protocols with bounded cheaters’ edge. *Proc. STOC’06*, pages 187–196, May 2006.
- [3] A. Caplin and B. Nalebuff. On 64%-majority rule. *Econometrica*, 56(4):787–814, July 1998.
- [4] J. S. Coggins and C. F. Perali. 64% majority rule in ducal Venice: Voting for the Doge. *Public Choice*, 97:709–723, 1998.
- [5] B. Cretney. The Election Methods Resource. <http://concordet.org/emr>, 2007.
- [6] H. Garcia-Molina. Elections in a distributed computing system. *IEEE Transactions on Computers*, C-31(1):49–59, January 1982.
- [7] M. Haahr. random.org site. <http://www.random.org>, 1998–2007.
- [8] J. Kurtzman and L. M. Koldau. Trombe, trombe d’argento, trombe squarciate, tromboni, and pifferi in Venetian processions and ceremonies of the sixteenth and seventeenth centuries. *Journal of Seventeenth-Century Music*, 5(1), 2002.
- [9] L. Lamport. Paxos made simple. *ACM SIGACT News*, 32(4):18–25, December 2001.
- [10] D. Lee. Living a doge’s life. *London Evening Standard*, October 2004.
- [11] M. Lines. Approval voting and strategy analysis: A venetian example. *Theory and Decision*, 20:155–172, 1986.
- [12] G. Maranini. *La Costituzione di Venezia*, volume 2. La Nuova Italia Editrice, 1931.

- [13] P. Murray. A distributed state monitoring service for adaptive application management. *Proc. International Conference on Dependable Systems and Networks (DSN'05)*, pages 200–205, 2005.
- [14] L. C. Noll, S. Cooper, and M. Pleasant. What is LavaRnd? <http://www.lavarnd.org/what/index.html>, 2001–2003.
- [15] J. J. Norwich. *A History of Venice*. Allen Lane, The Penguin Press, 1982.
- [16] D. Raines. Cooptazione, aggregazione e presenza al Maggior Consiglio: le casate del patriziato veneziano, 1297–1797. *Storia di Venezia - Rivista*, 1:1–64, 2003.
- [17] R. Richie and S. Hill. The case for Proportional Representation. *Social Policy*, 1996.
- [18] B. Schneier. *Beyond Fear*. Copernicus Books, Springer-Verlag New York Inc., 2003.
- [19] E. M. Tappan. *The World's Story: A History of the World in Story, Song and Art*, volume 5. Boston Houghton Mifflin, 1914.
- [20] J. Walker. HotBits site. <http://www.fourmilab.ch/hotbits/>, 1996–2006.

Appendix: Standard criteria for judging election protocols

Criterion A.1. (Majority) *If a candidate is the favourite of a majority of voters, that candidate must win.*

Criterion A.2. (Consistency) *For any way the voters are divided into two groups, if X is the winner for elections by both groups independently, X must also be the winner if the voters are not separated.*

Criterion A.3. (Pareto) *If candidate X is preferred to candidate Y by every voter, Y must lose.*

Criterion A.4. (Secret Preferences) *If candidate X wins, and some of the voters change their minds about the relative preferences they give to candidates to which they prefer X , X must still win.*

Criterion A.5. (Concordet) *If a candidate X pairwise beats every other candidate, then X must win the election.*

Criterion A.6. (Concordet Loser) *If a candidate X pairwise loses to every other candidate, X must lose the election.*

Criterion A.7. (Independence of Clones) *A clone set is a set S of candidates such that if any voter ranks a candidate Y higher than some Y_1 in S and lower than some Y_2 in S , then Y is in S . If S is a clone set of size ≥ 2 , and some candidate in S is eliminated and the election held again, then, if the original winner was in S , the winner for the new election must also be in S . If the original winner was not in*

S , then the original winner must also win the new election.

Criterion A.8. (Reversal Symmetry) *If candidate X wins (excluding ties), and all rankings by voters are reversed, then X must lose.*

Criterion A.9. (Smith) *The winner must be a member of the Smith set, which is the smallest nonempty set of candidates such for all Y in the set and Z not in the set, Y pairwise beats Z .*

Criterion A.10. (Local Independence from Irrelevant Alternatives) *If an election produces winner X , a new candidate Y is added and another election takes place, and Y is not in the Smith set, the new election must also produce winner X .*

Criterion A.11. (Monotonicity) *If candidate X loses, and then one or more voters change their minds so as to rank X in lower positions without changing the relative position of other candidates, then X must still lose.*

Criterion A.12 (Mutual Majority) *If there is a majority of voters for which it is true that they all rank a set of candidates above all others, then one of these candidates must win.*