Constitutional Political Economy (2020) 31:169–187
https://doi.org/10.1007/s10602-019-09290-6

ORIGINAL PAPER

How the Republic of Venice chose its Doge: lot-based elections and supermajority rule

M. Cristina Molinari[1]

Published online: 17 October 2019

Abstract

We study a family of voting rules inspired by the peculiar protocol used for over 500 years by the Republic of Venice to elect its Doge. Lot-based indirect elections have two main features: a pool of delegates is chosen by lot out of a general assembly, and then they vote in a single winner election with qualified majority. Under the assumption that the assembly is divided into two factions, we characterise the win probability of the minority and show that these features promote a more equitable allocation of political representation, striking a balance between protecting the minority and giving proper recognition to the majority. We then consider this family of voting procedures from a constitutional perspective: we analyse how the electoral result varies with the college size and the winning threshold in order to understand how these two parameters can be tuned when drawing up electoral law. We find that minorities are better off with larger majority thresholds. The role of the college size, on the other hand, is ambiguous: a smaller college size offers more protection to sparse minorities; for more sizeable ones, it depends instead on the qualified majority required for the election.

Keywords Voting · Minority protection · Probabilistic proportional representation · Sortition

JEL Classification D72 · N44 · C6

The author thanks Marco LiCalzi, Luciano Pezzolo, and two anonymous referees for useful comments. An earlier version of this paper circulated as WP 2018-2, Department of Economics, Ca' Foscari University of Venice.

✉ M. Cristina Molinari
cmolinar@unive.it

[1] Department of Economics, Ca' Foscari University of Venice, San Giobbe, Cannaregio 873, 30121 Venice, Italy

1 Introduction

In 1268, the aristocratic Republic of Venice introduced a bizzarre system to choose its *Doge*—its highest office as well as a position for life. The protocol, in force until the fall of the Republic in 1797, was an indirect election mechanism, because voting rights were restricted to an electoral college. The latter was chosen by a convoluted ten-round procedure, alternating voting with qualified majorities and sortition. Each intermediate voting round elected the subsequent nominating committee; only in the tenth and final round did the electoral college directly vote to elect the Doge. Across these rounds, the size of the college was alternately increased by voting and reduced by lot, as we now describe.

The active electorate lay within the Great Council (*Maggior Consiglio*), an assembly of male oligarchs. Its membership gradually increased from 480, when it was instituted in 1172, to more than two thousand (Norwich 1982: 278). The election protocol started by selecting at random thirty members of the Great Council, among those aged thirty and above. This first group of thirty was reduced by lot to a second committee of nine, which nominated a college of forty components, individually approved by a qualified majority of seven out of nine. This college of forty was reduced by lot to a committee of twelve, which nominated a college of twenty-five, individually approved by a qualified majority of nine out of twelve. This committee of twenty-five was reduced by lot to a college of nine, which elected a committee of forty-five, individually approved by a qualified majority of seven out of nine. A final round of sortition of the forty-five created a committee of eleven which chose, with a qualified majority of nine, the college of forty-one electors who actually voted to elect the Doge, with a *quorum* of twenty-five votes. In addition, no family was permitted to hold more than one member on each committee and the nominees' relatives were forbidden to vote (Norwich 1982: 198–200; Tucci 1982). Table 1 summarises the steps of the Venetian procedure for the dogal election.

This procedure seems designed to prevent a family or a coterie from planting its own candidate and, in so doing, gaining permanent control of the Republic. Other additional measures were taken to deter vote rigging, and prevent one faction

Table 1 The Venetian protocol for electing the Doge

Round	College size and minimum approvals
1	30 chosen by lot from the Great Council
2	9 chosen by lot
3	40 elected with the approval of 7/9
4	12 chosen by lot
5	25 elected with the approval of 9/12
6	9 chosen by lot
7	45 elected with the approval of 7/9
8	11 chosen by lot
9	41 elected with the approval of 9/11
10	1 elected with the approval of 25/41

dominating the electoral procedure through fraud or bribes. Voters were sequestered in a wing of the ducal palace until the Doge was elected, and were not allowed to communicate with anyone from outside. Windows were boarded over, a kitchen was set up, and beds were provided on the premises. Campaigning was entirely forbidden and punished as a crime against the state (da Mosto 1966: XIX–XXV). These efforts were not enough to avoid electoral intrigues entirely, but serious scandals were rare and they occurred mainly in the years of the Republic's decline.

Notwithstanding the elaborate procedure, elections were generally quite swift. On average, it took three days to complete the first nine steps and select the forty-one electors of the Doge. In the tenth and final round, the time needed to reach a majority of at least twenty-five votes was quite variable: sometimes a few hours were enough, but in 1618 the longest conclave took 30 days to elect Nicolò Donà (who, in an untimely manner, died 35 days later). Averaged over all the elections, the entire procedure was completed in 4 days in the fourteenth century, 3.25 days in the fifteenth century, 5.5 in the sixteenth century and 12 in the first half of the seventeenth century (da Mosto 1966: XVIII; Tucci 1982: 100).

This complex voting mechanism was the culmination of a series of changes in the electoral procedure. The first reform, enacted after the assassination of Doge Vitale II Michiel in 1172, introduced the indirect election: the right to vote, which was previously held by the popular assembly known as the *Concio,* was conferred on an electoral college of eleven persons appointed by the Great Council. In 1178, supermajorities and rounds were added; in the new system, the Great Council nominated four people who voted with a qualified majority of three out of four to designate forty people who then elected the doge by a simple majority. After 1229, the size of the final electoral college was increased to forty-one to prevent a tie (da Mosto 1966: XV–XVII; Maranini 1927: 173–78; Norwich 1982: 176–190). Reforms of the electoral procedure came to an end in 1268, when sortition was introduced in the selection of the forty-one electors interlaced with voting as described above. It took almost one hundred years of trial and error to fine-tune the electoral system, but once it was finally shaped it remained in force until the end of the Republic, more than five hundred years later.

The peculiarity of the Venetian electoral mechanism and its persistence over time, raise many questions: what was the purpose of an indirect election where the electoral college was chosen by lot? What were the consequences of imposing large supermajorities? Why was the procedure iterated so many times, varying the size of successive electoral colleges and the *quorum* required to decide?

The third question has already been partially addressed in the literature: Walsh and Xia (2011) show that using multiple iterations curbs manipulation of the voting process; Mowbray and Gollmann (2007) investigate how the electoral results of the Venetian protocol change with the number of iterations. This paper sheds light on the other two questions by studying a family of two-round voting mechanisms, called *lot-based indirect elections*, in which delegates chosen by lot out of a general assembly elect a single winner by qualified majority.[1] In other words, we examine

[1] In 1273, Venice adopted a two-round protocol to elect magistrates. An initial group of forty, chosen at random from the Great Council, was reduced by lot to a committee of nine people who elected each magistrate with a qualified majority of six out of nine. The use of iterations was later extended to magistrates'

an electoral procedure in which a single sortition step is followed by a single voting step. The extension to its iterated version is discussed at the end of the paper.

A lot-based indirect election combines the use of the lot and the requirement of a supermajority to reach a decision.[2] There is a body of literature on the protective role played by these two elements towards minorities; see Buchanan and Tullock (1962), Mueller et al. (1972), Mulgan (1984), and Schwartzberg (2014). However, as far as we are aware, there is no mathematical characterisation of how they offer such protection and, in particular, of the differences between the two instruments. This paper tries to fill this gap: we model lot-based indirect elections and study how their results are affected by two parameters, the size of the electoral college and the supermajority threshold.

We assume that the electorate is divided into two groups. This special case deserves attention because the Venetian aristocracy was divided into two main rival factions, the *Longhi* or *Case vecchie* (Long or Old houses) and the *Curti* or *Case nuove* (Short or New houses)—the old and the new nobility. These two clans often battled over the dogeship. The cohesion of the *Longhi* went beyond the dogal election, as they cooperated over a long period of time to achieve some control over the allocation of other official positions; in the Great Council the *Longhi* often voted as a bloc. (Norwich 1982: 414; Finlay 1980: 92–93 and 145–147).

With an electorate divided into two factions, the protection accorded to minorities by the voting protocol can be assessed as the probability that the winner is a candidate supported by the minority. We investigate the properties of this *minoritarian win probability* for different electoral protocols, and how it changes with the strength of the minority representation in the general assembly.

When qualified majorities and college selection by lot are used separately, we find that the supermajority rule is less favourable towards small minorities than sortition, but this conclusion is reversed for minorities with a sufficiently large membership. Thus, sortition not only makes the system less susceptible to fraud, which is a plausible reason why the Venetians introduced it (and kept it for over 500 years), but also affects how political representation is allocated.

We then study lot-based indirect elections from a constitutional perspective to understand the effects of the college size and the winning threshold; we do so carrying out a comparative statics exercise on how the minoritarian win probability changes with the number of delegates and with the inclusiveness of the rule for the college to reach a decision. We show that a higher threshold for a qualified majority increases the political representation of the minorities. The role of the college size is less clear-cut and depends on the supermajority threshold used: a large college

Footnote 1 (continued)
elections (Tucci 1982: 91). The usage of two consecutive uniform random draws in both the magisterial and the dogal election protocols is commonly viewed as a practical solution to the problem of ensuring a genuine random lot over a constituency as large as the Great Council.

[2] One of the earliest uses of a supermajority rule other than unanimity was the election of the Pope in 1179; selection by lot instead dates back to ancient Greece. Both were widely used in the Middle Ages. For a history of the two institutions see Chapter 3 of Schwartzberg (2014), Tridimas (2012) and Wolfson (1899).

size is bad news for a sparse minority, the more so the lower is the majority threshold; however, this is not necessarily the case for a larger minority. In other words, a smaller college size offers more protection only to sparse minorities.

Across all the electoral protocols considered in this paper, the minoritarian win probability is bounded above by the relative share that the minority has in the general assembly. This upper bound is reached under the unanimity rule or, for lower qualified majorities, when the electoral college is a singleton. This suggests that the supermajority threshold and the college size are substitutes with respect to the protection of the minority.

Our analysis shows that sortition and supermajorities, by offering some protection to the minority, could prevent any single house or clan from seizing power and, thus, could shield the Republic from the risk of becoming a hereditary monarchy.

A small body of literature has investigated the mathematical properties of the election protocols in the Republic of Venice. Lines (1986) focuses on the last election round and on approval voting. She compares the minority voters' incentive to misrepresent their preferences under either plurality or approval voting, and shows that the latter induces minorities to express their true preferences. Coggins and Perali (1998) look at the use of supermajority rule and, leveraging a result by Caplin and Nalebuff (1988), show that under the assumption of social consensus and for a large number of voters, the degree of supermajority used in the Venetian protocol produces a unique winner and avoids voting cycles. Our paper is most closely related to Mowbray and Gollmann (2007), who assume two factions and computationally simulate the original Venetian protocol, comparing it to some of its truncated variants. Their main result is that iterating the procedure protects the minority because it raises the minoritarian win probability. Finally, Walsh and Xia (2011) study the computational complexity of vote manipulation in a two-round lot-based election with many factions.

The remainder of the paper is organised as follows. Section 2 presents the model and introduces the minoritarian win probability. Section 3 analyses qualified majorities and sortition separately, comparing direct elections based on a supermajority rule to lot-based indirect elections where simple majority rule is used. Section 4 studies and elucidates the role played by the college size and the winning threshold in the protection of minorities. Section 5 concludes. All proofs are relegated to the appendix.

2 A model of lot-based indirect elections

We begin with presenting our setup for a lot-based indirect election with supermajorities. A general assembly of n individuals is to select a chair. Each of the n agents is a known member of one of two factions. Let $m \in (0, \frac{n}{2})$ denote the size of the minority faction and let $\mu = \frac{m}{n}$ be its percentage share of the total assembly size.

We assume that people from the same camp coordinate on a single candidate and focus on the electoral competition between the two factions, skirting the analysis of intra-party disagreement and negotiations. Given two undivided factions, their interaction ultimately boils down to a choice between their two preferred options. This simplifies the analysis of the electoral strategies because majority voting over two alternatives is non-manipulable: if a voter cannot gain from misrepresenting his preferences, we expect individuals to vote sincerely.[3]

An indirect election restricts the right to vote to an electoral college of size $c \leq n$, where n is the assembly size. Each *delegate* in the electoral college casts a vote for one of the two candidates. The election is won by the candidate who musters at least t votes out of the c available. To allow for supermajorities we let $t \in (\frac{c}{2}, c]$. For example, the simple majority rule corresponds to the winning threshold $t = \lfloor \frac{c}{2} \rfloor + 1$, where $\lfloor \cdot \rfloor$ is the floor function, and unanimity to $t = c$. For convenience, we use a roman t for the whole number of the winning threshold and a greek τ for its percentage value, with $\tau = \frac{t}{c} \in (\frac{1}{2}, 1]$.

Given a supermajority threshold, the election is undecided when neither candidate reaches t votes. If this happens, the voters must reconvene to negotiate an agreement and either faction has the power to stall the election. Since in this model consensus solutions have no role, there is no obvious prediction for how a decision may be reached. It seems reasonable to conjecture that *ex-ante* each candidate stands a chance to be the winner, but those with stronger support are more likely to win. Accordingly, we assume that a faction imposes its candidate with a probability proportional to its strength in the committee: when the ballot is indecisive, a faction with x members in an electoral college of size c wins with probability $\frac{x}{c}$.

Accordingly, the outcome of the election is completely determined by the composition of the college: delegates vote sincerely and the total tally for a candidate is equal to the number of delegates from his supporting faction who sit in the electoral college. As this number varies, three outcomes can occur: (a) if a faction controls at least t votes in the committee, then it can secure the election of its candidate; (b) if it has less than $c - t$ electors, then its candidate is a sure loser; and (c) when neither group controls enough votes to clinch the election, each of the two candidates has a positive probability of being the winner. Depending on the case, we say that a faction is (a) *decisive*, (b) *irrelevant* or (c) *influential*.

To model the use of sortition when forming the electoral college, we assume that the c delegates are chosen at random out of the n members from the general assembly. Since draws are without replacement, the number of minority faction's members in the electoral college is a random variable that follows the hypergeometric distribution.

Let $p_x[n, m, c]$ denote the probability that exactly $x = 0, 1, \ldots, c$ members from the minority faction sit in the college. The value of this probability depends on three parameters: the total number n of members in the assembly, the tally m of the minority faction and the size c of the electoral college.

[3] When choosing over two alternatives, the approval voting used in the Venetian electoral protocol is equivalent to the case where each voter only casts one ballot.

Using the binomial coefficient $\binom{n}{k}$ to denote the number of possible unordered combinations of n items, taken k at a time, we have

$$p_x[n,m,c] = \begin{cases} \dfrac{\binom{m}{x}\binom{n-m}{c-x}}{\binom{n}{c}} & \text{if}\, x \le m \\ 0 & \text{if}\, x > m \end{cases}$$

When convenient, we drop the square brackets and write p_x.

Before the electoral college is drawn, the outcome of a lot-based indirect election is described by the probability of winning of either faction. We focus on the probability $w[n, m, c, t]$ that the winner is the minority's candidate, called *minoritarian win probability* (MWP), and we compute it by taking the expected value of the electoral outcome over all possible college compositions. Given the probability distribution $p_x[n, m, c]$ over the number x of minority's delegates, we get

$$w[n,m,c,t] = \sum_{x=c-t+1}^{t-1} p_x[n,m,c]\left(\frac{x}{c}\right) + \sum_{x=t}^{c} p_x[n,m,c]. \tag{1}$$

The first sum refers to the electoral colleges in which the faction is influential and thus has a probability of winning proportional to its share $\frac{x}{c}$ of the college; the second sum refers to the colleges in which it is decisive and therefore it elects its candidate. As for p_x, we drop the parameters in square brackets and write w when convenient.

The probability w is a measure of the protection that the electoral protocol gives to the minority faction. The higher its value, the more chances the minority has to be the winner of the election. In what follows we investigate the properties of w under alternative versions of the electoral protocol.

3 Supermajorities, sortition and minorities

A lot-based indirect election with supermajorities is a voting protocol that mixes two ingredients: (a) the use of the sortition and (b) the requisite of a large consensus to reach a decision. In this section we give a formal characterisation of the effects that chance and qualified majorities play, with respect to the political representation of minorities. Our purpose is to unveil the differences between these two instruments. To this end, we compare two protocols: (1) a direct election without sortition in which everyone votes and decisions are reached with a qualified majority rule; and (2) a voting protocol in which the electoral college is selected at random but decisions are taken under simple majority rule. For both protocols we characterise the probability w that a candidate backed by less than half the electorate end up being the winner.

3.1 Minority protection in direct elections with supermajorities

Consider an election in which all the n individuals have the right to vote so that $c = n$. Because the number of minority individuals voting is $m < n/2$, a minority is never decisive, regardless of the majority threshold t. This implies $w < 1$ in any non-randomized electoral procedure. Any political representation the minority might receive is associated with it being influential and, therefore, it is at most equal to its percentage share in the general assembly. Since in a direct election a faction is influential when $n - t < m < t$ we conclude that, for any supermajority threshold $t \in (\frac{n}{2}, 1]$, the MWP is given by:

$$w = \begin{cases} 0 & \text{if}\, m < n - t \\ \dfrac{m}{n} & \text{if}\, m \geq n - t \end{cases}$$

Dividing the right-hand side by n and using $\mu = \frac{m}{n}$, this can be conveniently rewritten as:

$$w = \begin{cases} 0 & \text{if}\, \mu < 1 - \tau \\ \mu & \text{if}\, \mu \geq 1 - \tau \end{cases}$$

To compare the minority protection across protocols, we give a graphical representation of the MWP as a function of the minority percentage allegiance μ. Figure 1 plots w for three possible values of the winning threshold: simple majority rule, a supermajority of four out of five and the unanimity rule. The horizontal axis is for the fraction μ of minority members in the general assembly and the vertical axis for the MWP. Clearly, in our discrete setting w is a collection of dots corresponding to values of $\mu = \frac{1}{n}, \frac{2}{n}, \dots$, but, as a matter of convenience, here and in the following figures, we plot it as a continuous function. As shown in the picture, the effect of the supermajority is to allocate some "power" to minority factions that are not too small but that would otherwise be cut out from the decision process (see the leftmost panel of Fig. 1). As the decision rule is made more inclusive, i.e. for a larger winning threshold t, the protection increases, as we can see moving rightward in Fig. 1.

The degree of protection reaches its maximum when unanimity is required. In the latter case the minority has a probability of winning equal to its percentage representation in the assembly. We can interpret this situation as a probabilistic version of proportional representation.[4] Given that proportional representation maximises

[4] In a single candidate election, *ex-post* proportional representation is impossible, except in the case of unanimity. However, it can be attained *ex-ante* when the probability of winning the seat reflects the distribution of votes. For a discussion of probabilistic proportional representation, see Amar (1984).

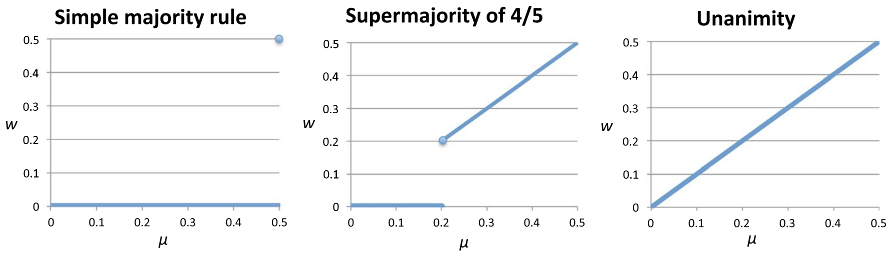

Fig. 1 Minoritarian win probability in a direct election

political equality, supermajoritarian rules in single district elections can be interpreted as a means to treat factions more equally.

A second effect of larger supermajority thresholds is that negotiations to secure a winning majority occur more often, leading the parties to strive for more inclusive decision making. This can be seen in Fig. 1: whenever the probability line lies on the diagonal $w = \mu$, the winner is chosen through negotiations. Thus, supermajority provisions also engage majority and minority into bargaining. This can be a positive aspect of the procedure as long as the search for a deal can be handled; if, instead, the underlying culture of the political system is characterised by strong interclan antagonism, a system that often requires seeking a consensual solution could lead to serious political strife, including civil war and the emergence of alternative institutions. This possibility is well illustrated by the markedly different experiences of the two, otherwise quite similar, maritime republics of Genoa and Venice. Genoa was characterised by aristocratic families organised in a rigid clan structure, the *Alberghi*, that hindered cooperation and, more than once in its history, it had to resort to some external authority to attain political stability. It also underwent several constitutional reforms. Venice, instead, was more successful in securing wide consensus and political order (Greif and Laitin 2004).

The last feature of a supermajoritarian election procedure we want to stress is related to its manipulability. Except for the two extreme cases of simple majority and unanimity, the MWP is discontinuous. Going beyond a simple model where voters cannot change their party allegiance, such discontinuities provide strong incentives for bribery, because a small number of turncoat voters may lead to a markedly different result. This suggests that a discontinuity in the probability of winning may be an undesirable feature of the election protocol. The inclusion of all sorts of precautions to avoid vote manipulation was standard in the statutes of medieval Italian city states (Wolfson 1899). This confirms that fraud in voting was perceived as an extremely dangerous practice for the independence of the city state, and that any effort to avoid it had to be made. As we show in the following section, the issue of vote manipulation is less compelling in lot-based elections.

3.2 Minority protection in lot-based indirect elections

We now turn our attention to elections where the set of voters is reduced by lot. As above, the outcome of the election depends only on the composition of the electoral committee but selection by ballot turns the number of minority delegates into a random variable. In particular, whereas in a direct election a minority is never decisive, when voters are reduced by lot a minority has a chance to be "upgraded" and become a majority within the electoral college. When $m > c$, there is a positive probability the minority becomes decisive and can impose its candidate. Nonetheless, this probability is smaller than the probability that a minority in the general assembly remains a minority in the electoral college. We formalise this intuition in Proposition 1 which will later be used to study the degree of protection of different lot-based voting protocols.

Proposition 1 *If* $m < n/2$, $p_x[n, m, c] \geq p_y[n, m, c]$ *for any* $x < y$ *such that* $x + y = c$. *Moreover* $p_x[n, m, c] > p_y[n, m, c]$ *if* $x \leq m$.

The following two propositions characterise the minoritarian win probability by imposing a lower and a upper bound on w.

Proposition 2 *If* $m \geq c - t + 1$, *the* MWP *in a lot-based indirect election is strictly positive.*

Proposition 2 clarifies that the selection by lot of the electoral college gives some representation to the minority faction unless the latter is too small. Given the size c of the electoral college and the winning threshold t, the minority has a positive probability to see its candidate winning unless it has not enough members in the general assembly to hope to contend a sure majority to the rival faction in the college. This result confirms that the use of sortition in election protocols promotes a more equitable allocation of political representation.

A representation benefit for the minority, however, is a representation deficit for the majority. Is such deficit too harsh? Proposition 3 shows that the representation afforded by the protocol to minorities is not "unnecessarily generous" because it is always smaller that the relative share μ that the minority has in the general assembly or, in our probabilistic interpretation, that the minority faction never gets more than proportional representation.

Proposition 3 *For any value of* n, m, c *and* $t \in (c/2, 1]$, *the minority winning probability* w *in a lot-based indirect election is bounded above by* μ. *In particular,* $w = \mu$ *under the unanimity rule.*

To illustrate how sortition affects the minority's chance of winning, in Fig. 2 we plot the MWP for $n = 100$ and $c = 10$, under a simple majority rule. As in Fig. 1, the horizontal axis shows the fraction μ of minority members in the general assembly and the vertical axis the probability w that the minority wins; the dotted bisector $w = \mu$ represents proportional representation. As illustrated in Fig. 2, the MWP is bounded away from zero for any non-trivial minority and it is bounded above by proportional representation.

Compared to Fig. 1, Fig. 2 shows that both supermajorities and sortition are ways to protect the minorities. In both cases the representation gained by the minority is smaller than proportional representation, but sortition has the effect of smoothing out the changes in representation associated with changes in the strength of the minority: the probability of winning in lot-based indirect elections moves smoothly with μ and this, as discussed before, can possibly reduce the incentive to manipulate the election by inducing some voters to switch allegiance.

4 The role of winning thresholds and college size

The adoption of a lot-based indirect election protocol involves the choice of two key parameters: the majority threshold t and the college size c. To evaluate this constitutional choice, we turn to a comparative statics analysis of the role of t and c.

4.1 Minoritarian win probability and supermajorities

We fix the number n of people in the general assembly and the membership m of the minority faction. We analyse the effect of changing the majority threshold t for a given size c of the electoral college, by focusing on how the majority threshold of a lot-based indirect election determines the MWP. The value of t decides whether a faction is influential, decisive or irrelevant in the electoral college, but it does not affect the college composition nor its probability distribution. In other words, if we

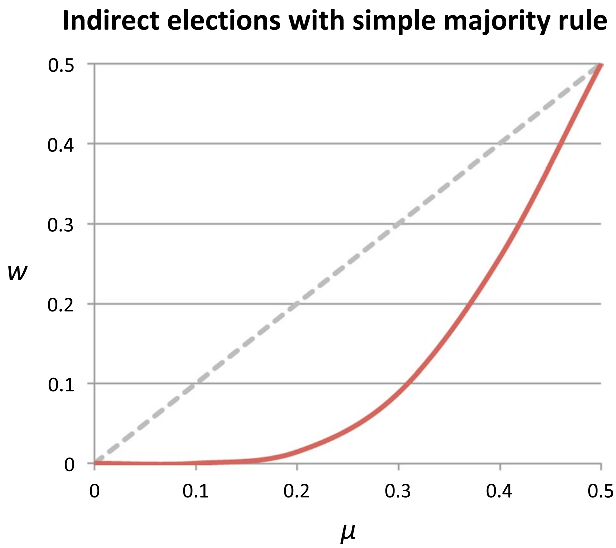

Fig. 2 Minoritarian win probability ($n = 100, c = 10, \tau = 0.5$)

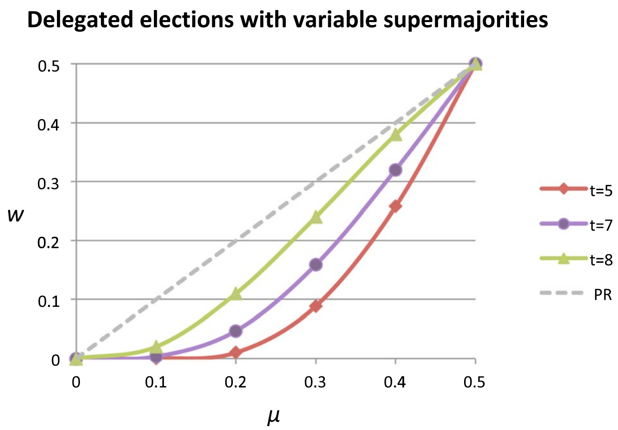

Fig. 3 Minoritarian win probability ($n = 100, c = 10$)

look at Eq. (1), the supermajority threshold t pins down which values of x enter in the first summation (for influential colleges), which enter the second summation (for decisive colleges), and which are left out (for irrelevant ones).

When t is increased, there are fewer irrelevant and decisive colleges, whereas the number of influential colleges gets larger. How is MWP affected by these changes? There is a positive effect because the minority is less often irrelevant; and there is a negative effect because it is less likely to be decisive. The next proposition states that the first effect is stronger than the second: overall, a larger winning threshold increases the MWP.

Proposition 4 *In a lot-based indirect election, the MWP increases with the supermajority threshold* t, *and it is equal to proportional representation under the unanimity rule.*

The intuition behind Proposition 4 is straightforward. Raising the supermajority threshold has both positive and negative consequences: the positive effect is associated to fewer instances in which a faction is irrelevant and the negative one with fewer instances in which it is able to control the result of the election. In absolute terms, gains and losses are equivalent; however, if we weigh them for the probabilities to occur, the equivalence fails. In fact, the positive consequences accrue in situations where the faction has few delegates in the electoral college, whereas the negative impact is produced when it has many. And, by Proposition 1, a minority is more likely to end up with a small representation in the electoral college than with a large one. Therefore, taking into account the odds, the positive effect prevails.

The effect of raising the supermajority threshold is illustrated in Fig. 3 where the probability of winning is again computed for $n = 100$ and $c = 10$. The three curves correspond to different values for t, namely $t = 5$ (diamonds), $t = 7$ (circles) and $t = 8$ (triangles). As per Proposition 4, raising t makes the MWP go up and approach the dotted line that corresponds to proportional representation. Notice that the MWP

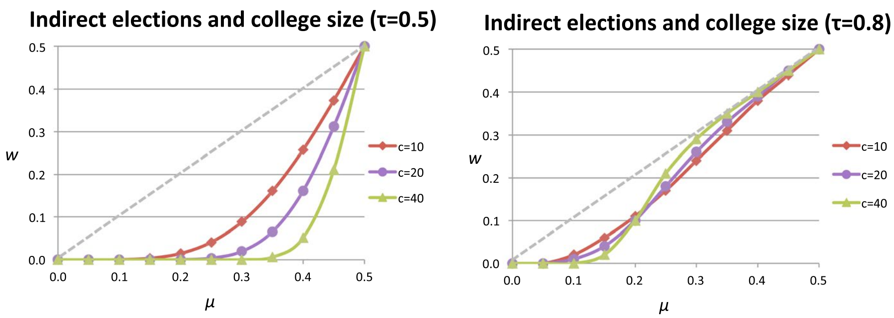

Fig. 4 Minoritarian win probability ($n = 100$)

is nil when $m \leq c - t$ and strictly positive otherwise. This implies that negligible minorities are less often excluded from participation for larger t's and, therefore, the curves raise above the horizontal axis for μ closer to zero. We conclude that, in indirect elections, the supermajority rule reinforces the minority protection effect associated with college selection by lot.

4.2 Minoritarian win probability and college size

We now turn our attention to the size of the electoral college. A comparative statics exercise on the parameter c requires some caution. In fact, if the supermajority threshold t is given, an increase in the college size loosens the majority requirement. Therefore, to single out the effect of the college size, we fix the percentage supermajority threshold $\tau = \frac{t}{c}$, and let t rise with c. Notice that, when τ is fixed, few values of c correspond to a well-defined model: our election problem is discrete, so it is meaningful only if both c and t are integers. In the following analysis we leave it implicit that changes in c correspond to an integer value of $t = \tau c$.

The comparative statics on the college size is less straightforward than for the majority threshold of Sect. 4.1 because changes in c affect not only the electoral outcome, but also the composition of the electoral committee itself and its probability distribution. For this reason, we give general results only for extreme values of the college size. The behaviour of the MWP for intermediate values of c is discussed on some examples.

Consider first the extreme cases. At the lowest possible size for the electoral college ($c = 1$), the unique delegate is decisive. This implies that the MWP for a faction of m members equals the probability of having one of its member chosen to sit in the college, out of a total population of n individuals. This yields $w = \frac{m}{n} = \mu$. Therefore, the smallest college size corresponds to proportional representation. At the opposite extreme, the electoral college can have a size $c = n$. This is the case of direct elections analysed in Sect. 3—and represented in Fig. 1—where we showed that a minority faction enjoys proportional representation only if its allegiance is greater than the supermajority threshold; otherwise, it is irrelevant and receives no representation.

To illustrate what happens for intermediate values of the college size, we compute the MWP using the closed-form formula given in (1). The results are summarised by the two prototypical situations shown in Fig. 4, where we plot the MWP for $n = 100$ and three different values of the college size: $c = 10$ (diamonds), $c = 20$ (circles) and $c = 40$ (triangles). The percentage supermajority threshold is $\tau = 0.5$ in the left panel and $\tau = 0.8$ in the right one.

As Fig. 4 shows, the impact of having a larger college depends on the threshold τ. To understand why, note that a small minority party might run short of candidates to fill a large electoral college. In particular, if $m \leq c - t$, the minority becomes irrelevant regardless of its luck in the sortition. This is equivalent to $m \leq c(1 - \tau)$ or, if we divide through by n, to

$$\mu \leq \frac{c}{n}(1 - \tau).$$

This inequality shows that, when the electoral college size increases, the upper bound on μ that deprives a faction of political representation gets larger for any value of τ, but the increase is larger when the supermajority threshold is small, and it becomes nil for $\tau = 1$.[5] This can be observed in Fig. 4 where, as c increases, the minoritarian win probability w shrinks to zero for low values of μ, but it does so more noticeably when τ is smaller.

For larger values of μ, the difference in the two panels of Fig. 4 can be traced back to the first two panels of Fig. 1, representing the MWP when $c = n$, for $\tau = 0.5$ and $\tau = 0.8$ respectively. As the college size grows and approaches n, the minority obtains a result close to what it would get in a direct election. Therefore, the winning probability shrinks to zero as long as $\mu < \frac{1}{2}$ when $\tau = 0.5$, but it converges to proportional representation for large values of μ, when $\tau = 0.8$. This is why the MWP is s-shaped for the large majority threshold $\tau = 0.8$.

Altogether, increases in the college size are associated with a lower representation for a minority with relatively few members (μ close to 0). The effect for a more sizeable minority is ambiguous and it depends on the supermajority threshold.

5 Conclusions

We analyse a class of indirect election mechanisms where the electoral college is chosen at random out of a general assembly and a supermajority is required to secure the winner. Our interest was spurred by a related mechanism in use in the patrician Republic of Venice for more than 500 years.

Under the restriction that there are only two factions, we study the minority faction's probability of winning. We show that this mechanism, compared to a direct election, protects the minority by guaranteeing a higher probability of winning, without unduly penalising the majority, because the probability always remains

[5] From Proposition 3 we already know that with the unanimity rule ($\tau = 1$) we get proportional representation for any college size.

smaller than the proportion of the minority faction in the assembly. In other words, lot-based indirect elections strike a balance between protecting the minority and giving proper recognition to the majority.

We investigate how the protection of the minority varies with the college size and the supermajority threshold. We find that a smaller college size is more protective only for small minorities and that larger supermajority thresholds are more protective.

The family of election procedures studied in this paper departs from the Venetian protocol in two respects: (a) we restrict our attention to two-round procedures, and (b) we look at the case of an electorate divided into just two factions. Extensions to multi-round procedures can be done along the lines of Mowbray and Gollmann (2007); their results suggest that iterating selection by lot and qualified majority elections many times raises the minoritarian win probability. A more interesting extension would go beyond (b) and study lot-based indirect elections with more than two factions. This would require an analysis of strategic voting on the one hand, and of coalition formation on the other.

Appendix

Proof of Proposition 1

Proof First, consider $x \leq m$, which implies $p_x > 0$. Then either $y > m$, in which case $p_y = 0$ and thus $p_x > p_y = 0$, or else $y \leq m$ and $p_y \neq 0$. In the latter case the definition of the binomial coefficient and $y = c - x$ imply:

$$\binom{m}{x} = \frac{m!}{x!(m-x)!} = \binom{m}{y}\frac{y!(m-y)!}{x!(m-x)!} = \binom{m}{y}\frac{(c-x)!(m-c+x)!}{x!(m-x)!}. \tag{2}$$

By an analogous reasoning we obtain:

$$\begin{aligned}\binom{n-m}{c-x} &= \binom{n-m}{c-y}\frac{(c-y)!(n-m-c+y)!}{(c-x)!(n-m-c-x)!}\\ &= \binom{n-m}{c-y}\frac{x!(n-m-x)!}{(c-x)!(n-m-x-c)!}.\end{aligned} \tag{3}$$

We now use (2) and (3) to express p_x as a function of p_y:

$$
\begin{aligned}
p_x &= \frac{\binom{m}{x}\binom{n-m}{c-x}}{\binom{n}{c}} \\
&= \frac{\binom{m}{y}\binom{n-m}{c-y}}{\binom{n}{c}} \frac{(m-x-c)!}{(m-x)!} \frac{(n-m-x)!}{(n-m-x-c)!} \\
&= p_y \frac{(m-x-c)!}{(m-x)!} \frac{(n-m-x)!}{(n-m-x-c)!} \\
&= p_y \frac{(n-m-x)}{m-x} \cdot \frac{(n-m-x-1)}{(m-x-1)} \dots \frac{(n-m-x-c+1)}{(m-x-c+1)} > p_y,
\end{aligned}
$$

where the last inequality follows from the fact that, being $m < n/2$, each fraction in the last line is greater than one.

Next, consider the case $x > m$. Since $x < y$, we have $y > m$ and, therefore, $p_x = p_y = 0$. This completes the proof. □

Proof of Proposition 2

Proof Suppose $w = 0$. Given the definition of w in (1), it must be $p_x = 0$ for any $x \geq c - t + 1$. But with the hypergeometric distribution $p_x = 0$ if and only if $x > m$. It follows that $m < c - t + 1$, a contradiction. □

Proof of Proposition 3

Proof First we prove that $w \leq \mu$. Start with Eq. (1):

$$w = \sum_{x=c-t+1}^{t-1} p_x \frac{x}{c} + \sum_{x=t}^{c} p_x.$$

Add and subtract both $\sum_{x=0}^{c-t} p_x \frac{x}{c}$ and $\sum_{x=t}^{c} p_x \frac{x}{c}$ to get

$$w = \sum_{x=0}^{c} p_x \frac{x}{c} - \sum_{x=0}^{c-t} p_x \frac{x}{c} - \sum_{x=t}^{c} p_x \frac{x}{c} + \sum_{x=t}^{c} p_x.$$

For the hypergeometric distribution the expected value $\sum_{x=0}^{c} p_x x$ is equal to $\frac{m}{n}c$. Inserting this value in the equation above and rearranging the terms yields

$$w = \frac{m}{n} - \sum_{x=0}^{c-t} p_x \frac{x}{c} + \sum_{x=t}^{c} p_x \left(1 - \frac{x}{c}\right)$$

We can simplify this expression and switch to a new summation index in the last term, yielding:

$$\begin{aligned} w &= \frac{m}{n} - \sum_{x=0}^{c-t} p_x \frac{x}{c} + \sum_{x=0}^{c-t} p_{c-x} \left(1 - \frac{c-x}{c}\right) \\ &= \frac{m}{n} - \sum_{x=0}^{c-t} (p_x - p_{c-x}) \frac{x}{c}. \end{aligned} \tag{4}$$

We now use Proposition 1 to sign the last term of this equality. By definition of supermajority threshold, $t > \frac{c}{2}$. Then, for any value of $x \in \{0, 1, \dots, c-t\}$, we know that $x < c - x$. Since $x + (c - x) = c$ the proposition applies and, thus, each term $(p_x - p_{c-x})$ in the sum is non-negative. This implies $w \leq \mu$.

To complete the proof consider $t = c$. Then neither faction can be decisive and Eq. (4) reduces to

$$w = \frac{m}{n} - \sum_{x=0}^{0} (p_x - p_{c-x}) \frac{x}{c} = \frac{m}{n} = \mu.$$

□

Proof of Proposition 4

Proof Fix n, m and c. Abusing notation, we drop these three parameters and write $w[t]$ for the MWP given a value of the supermajority threshold $t < c$. It is enough to show that when we raise the threshold from t to $t + 1$, the difference $\Delta_t w = w[t+1] - w[t]$ is positive.

Using the definition of MWP given in Eq. (1) we write

$$w[t] = \sum_{x=c-t+1}^{t-1} p_x \frac{x}{c} + \sum_{x=t}^{c} p_x.$$

Similarly, the winning probability for a threshold $t + 1$ is

$$w[t+1] = \sum_{x=c-t}^{t} p_x \frac{x}{c} + \sum_{x=t+1}^{c} p_x.$$

Notice that all the terms in $w[t]$ and $w[t+1]$ are the same except for those indexed by $c-t$ and t. Then the difference $\Delta_t w$ reduces to:

$$\Delta_t w = p_{c-t}\frac{c-t}{c} + p_t\frac{t}{c} - p_t = \frac{c-t}{c}(p_{c-t} - p_t). \tag{5}$$

To determine the sign of $\Delta_t w$ notice that $c - t \geq 0$ by definition of supermajority threshold. This implies that $\Delta_t w$ has the same sign as $(p_{c-t} - p_t)$. Since the pair of probabilities p_{c-t} and p_t satisfies the conditions of Proposition 1 and we are computing w for a minority faction, we know that $(p_{c-t} - p_t) \geq 0$ and thus conclude that $\Delta_t w \geq 0$.

To complete the proof consider the unanimity rule. When $t = c$ the MWP is:

$$w = \sum_{x=1}^{c-1} p_x \frac{x}{c} + p_c = \frac{1}{c}\left(\sum_{x=1}^{c} p_x x - p_c c\right) + p_c = \frac{1}{c}\sum_{x=1}^{c} p_x x = \frac{m}{n}$$

where the last equality follows from the fact that the expected value $\sum_{x=1}^{c} p_x x$ is equal to $\frac{mc}{n}$ for a hypergeometric random variable. □

References

Amar, A. R. (1984). Choosing representatives by lottery voting. *Yale Law Journal*, *93*(7), 1283–1308.

Buchanan, J. M., & Tullock, G. (1962). *The calculus of consent*. Ann Arbor: University of Michigan Press.

Caplin, A., & Nalebuff, B. (1988). On 64%-majority rule. *Econometrica*, *56*(4), 787–814.

Coggins, J. S., & Perali, C. F. (1998). 64% majority rule in ducal Venice: Voting for the Doge. *Public Choice*, *97*(4), 709–723.

da Mosto, A. (1966). *I dogi di Venezia nella vita pubblica e privata* (2nd ed.). Milano: Aldo Martello Editore.

Finlay, R. (1980). *Politics in Renaissance Venice*. New Brunswick: Rutgers University Press.

Greif, A., & Laitin, D. D. (2004). A theory of endogenous institutional change. *American Political Science Review*, *98*(4), 633–652.

Lines, M. (1986). Approval voting and strategy analysis: A Venetian example. *Theory and Decision*, *20*(2), 155–172.

Maranini, G. (1927). La costituzione di Venezia. Dalle origini alla Serrata del Maggior Consiglio. [Reissued in 1974]. Firenze: La Nuova Italia.

Mowbray, M., & Gollmann, D. (2007). Electing the Doge of Venice: Analysis of a 13th century protocol. In *Proceedings of the 20th IEEE computer security foundations symposium* (pp. 295–310). IEEE Computer Society.

Mueller, D. C., Tollison, R. D., & Willett, T. D. (1972). Representative democracy via random selection. *Public Choice*, *12*(1), 57–68.

Mulgan, R. G. (1984). Lot as a democratic device of selection. *The Review of Politics*, *46*(4), 539–560.

Norwich, J. J. (1982). A history of Venice. [Reissued in 2003]. London: Penguin Books.

Schwartzberg, M. (2014). *Counting the many: The origins and limits of supermajority rule*. Cambridge: Cambridge University Press.
Tridimas, G. (2012). Constitutional choice in ancient Athens: The rationality of selection to office by lot. *Constitutional Political Economy*, *23*(1), 1–21.
Tucci, U. (1982). I meccanismi dell'elezione dogale, reissued in 2014 in Venezia e dintorni. Evoluzioni e trasformazioni. Roma: Viella.
Walsh, T., & Xia, L. (2011). Venetian elections and lot-based voting rules. In E. Elkind, U. Endriss & J. Lang (Eds.), *Proceedings of the IJCAI-2011 workshop on social choice and artificial intelligence* (pp. 93–98).
Wolfson, A. M. (1899). The ballot and other forms of voting in the Italian communes. *The American Historical Review*, *5*(1), 1–21.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.