Voting as a lottery

Published version:

DOI:10.1016/j.jpubeco.2016.11.016

Voting as a Lottery*

Giuseppe ATTANASI† Luca CORAZZINI‡ Francesco PASSARELLI§

November 2016

Abstract

This paper studies the issue of constitutional design, and supermajorities in particular, from a behavioral economics perspective. The relevant parameters are voting power, risk aversion, and pessimism. Voters who feel powerful prefer lower thresholds, while risk averters and those who feel pessimistic about the majority prefer higher thresholds. We also analyze the effects of loss aversion and overconfidence. The former leads voters to prefer more protective voting rules, a manifestation of their bias towards the status quo. The latter leads them to prefer overly low (high) protection when they receive good (bad) news about how others will vote. Finally, we study constitutional agreements on the voting rule. Members of the constituent assembly are heterogeneous in the parameters above. Weak and minority members anticipate high expropriation risk in future decisions. This gives them consistent leverage to push for a protective constitution.

Keywords: majority rule, supermajority, risk aversion, weighted votes, loss aversion, overconfidence, behavioral political economy, constitutions, tyranny of the majority.

*We thank Francesco Trebbi, two anonymous referees, Philippe Aghion, Olivier Armantier, John Carey, Dean Lacy, Matthias Messner, Aldo Montesano, Antonio Nicolò, Anastasiya Shchepetova, Robert Sugden, Piero Tedeschi, Unal Zenginobuz and participants at the PET 2008 in Seoul, the PET 2009 in Lyon and a seminar at Catholic University Milan, for useful comments and suggestions. We also thank Brandon Gill for outstanding assistance in editing. G. Attanasi gratefully acknowledges financial support by the project "Creative, Sustainable Economies and Societies" (CSES), University of Strasbourg IDEX Unistra. F. Passarelli gratefully acknowledges Collegio Carlo Alberto for financial support.

†BETA, University of Strasbourg. Email: attanasi@unistra.fr

‡Department of Economics (SEAM), University of Messina. Email: lcorazzini@unime.it

§Contact author: University of Turin and Bocconi University. Email: francesco.passarelli@unibocconi.it.

We treat voting as a lottery in which the probability of winning or losing depends on the majority threshold. The stylized situation is a legislature in which two reform proposals are put forward for voting. A voter "wins" if her most preferred proposal passes. She "loses" if the other one passes, as we assume that it is worse than the status quo. The outcome is uncertain because she does not know how the others will vote. In this lottery, the probability of any outcome depends on the majority threshold. If the threshold is low, decisions are easy to make. However, with a low threshold not only winning, but also losing are quite likely. There is a trade-off: decisiveness versus protection. Demand for decisiveness comes from three types of agents: the powerful ones (say powerful parties, large political factions, or big states in federal contexts), those who believe they share the same preferences with the majority of people, and those obtaining large gains from winning. Demand for protection comes from risk averters or weak minorities with much to lose from tyranny or expropriation.

There is no bargaining at the "legislative" stage in this model: the two competing proposals are exogenous. This is plausible when, say, two parties make divergent proposals that reflect the broadest consensus of their electoral base, and no side-payments are possible in order to collect a wider support. Society is split into two groups, and each proposal benefits only one group or the other. Voters ignore the exact size of the two groups. This type of ignorance is plausible for several reasons. For instance, voters may have limited information about others' preferences. The behavior of swing voters may be highly unpredictable. Pre-election polls may be subject to mistakes. Secret or last-minute agreements among factions may change the composition of each group.

Within this framework, we endogenize the preferred majority threshold of a voter. Under simple conditions, this threshold is unique and, interestingly, it is often either the bare majority or unanimity.[1]

Next, we use individual preferences for majority thresholds to build a constitutional stage. The members of the constituent assembly ignore future decisions, but they do know whether

[1] Here is the intuition for a corner solution. If for any threshold, winning is more likely than losing and the cost of losing is (weakly) lower than the benefit of winning, then an agent prefers the lowest possible threshold: the bare majority. If for any threshold, winning is less likely than losing, then the agent becomes a risk-seeker. Her optimization problem yields one of the two most risky values of the threshold, either simple majority or unanimity. She always chooses the latter, except when the relative advantage of winning is "sufficiently large" (cf. Proposition 1).

they are likely to belong to a minority or not. They also know whether they will have high or low influence on future voting. Minorities or weak members have more at stake, since they anticipate high expropriation risk in future decisions. This gives them consistent leverage to push for a high threshold. Our claim is that negotiations amongst "unequals" are likely to yield quite protective constitutions.

In reality, the simple majority threshold occurs less frequently than one may think. There are many examples of supermajorities that cross countries and industry lines. Most governments adopt bicameral systems, which are de facto supermajorities. The U.S. Federal Constitution requires a two-thirds majority to override a presidential veto, to ratify a treaty, or to expel a member of Congress. Three-fifths of the full Senate must approve any waiver of balanced budget provisions. Recently, the Lisbon Treaty has adopted a double supermajority for the EU Council. In international treaties, members can exercise vetoes when decisions concern their crucial interests (e.g. the Council of the EU, the UN Security Council). In corporate boards, high thresholds are usually required to pass major actions (e.g. mergers and acquisitions, major capital expansions).

Knight (2000) argues that pro-tax state legislatures tend to adopt supermajorities because the median legislator has an incentive to give up her pivotal role in order to reduce the agenda-setting power of extremists in her party. Dal Bo (2006) suggests that appropriate supermajority requirements induce the right conservative bias, solving possible time inconsistency problems in policy making. Simple majority is a better option when policies are time-consistent.

Traditionally, it has been argued that supermajorities can mitigate the "tyranny of the majority" (Buchanan and Tullock, 1962). Rae (1969) suggests that simple majority is the only rule that minimizes the expected cost of being a part of the minority. His approach has inspired several subsequent papers that extend the probabilistic voting model in different dimensions. Curtis (1972) considers a committee whose members may have heterogeneous probabilities to vote for any proposal. He proves that simple majority is the only rule that satisfies the utilitarian criterion. Badger (1972) uses the same model and shows that voters' preferences over majority thresholds are single-peaked, which in turn implies that, had the members to vote on the threshold, a Condorcet winner would always exist. Barberà and Jackson (2004) prove that members' bliss points can be ordered according to their (subjective) probability of voting for the reform.[2] Coelho (2005) uses a similar model to investigate the voting rule selected by the normative Rawlsian maximin criterion. All of these papers consider an agent who does not know today whether she will vote in favor or against the reform tomorrow. The reform is always posed against the status quo, and the relative benefit of the most preferred alternative is normalized to 1. Unlike these models, we consider an agent who is aware of her preferences, but unaware of how many other voters share her preferences. We explicitly consider the cost of being expropriated, which may be large for voters with extreme preferences or for members of minority groups. Our representative agent is also aware of her power, which we model with weighted votes. Higher weight leads voters to prefer a lower supermajority. This idea is certainly not new, but it has not been fully developed theoretically, to the best of our knowledge.

The present paper is also related to a relevant body of literature that studies voting equilibria when a voter's beliefs about the behavior of other voters depend on the distribution of voter types (e.g. Laslier, 2009; Myerson and Weber, 1993). This literature assumes that beliefs are endogenous, whereas voting rules are exogenous. We do the opposite: we endogenize voting rules while keeping beliefs exogenous. We also study how voters change their preferences for voting rules as a result of Bayesian updating of their beliefs.

Our model is designed to accommodate behavioral distortions into the process of constitutional design that are novel to the voluminous literature on this topic. We consider loss aversion and overconfidence. A loss averse individual wants more protection against the risk of being expropriated. Hence, she prefers less decisive voting rules. This preference is related to the status quo bias emphasized by Alesina and Passarelli (2015). Further, an overconfident individual exaggerates her reactions to new information. If a signal contains good news (i.e., the signal says that more people than she thought will vote like her), then she prefers overly decisive rules. This is the case considered by Ortoleva and Snowberg (2015) who claim that overprecise voters display extremeness in political behavior. In our model, this reaction to good news is "translated" into demand for more decisive rules, which leads to more radical changes. In the case of bad news, an interesting trade-off comes about. On the one hand, the agent wants more protection because the news is bad; on the other hand, she wants more decisiveness, because

[2] For a thorough welfarist assessment of majority thresholds, see Beisbart and Bovens (2007), and references therein.

her uncertainty is lessened after the signal. However, for sufficiently large overconfidence, there is no trade-off. Bad news always yields demand for more protective rules.

Other papers have addressed the same "constitutional" question as in this paper: "Which voting rule should/will a group adopt?" Aghion and Bolton (2002) show that the optimal majority threshold is increasing in the expected cost of compensating the losing minority, when agents do not know ex-ante if they will lose or gain from the provision of a public good. Messner and Polborn (2004) suggest that, relative to young voters, the older voters are more conservative and prefer higher thresholds because they pay more taxes and have less opportunities to benefit from fiscal returns. Thus aging societies adopt more conservative voting rules. Barberà and Jackson (2004) study the self-stability of voting rules. In the case of homogeneous committees, simple majority is always self-stable, while in the case of a heterogeneous committee self-stability is not guaranteed. In a related paper, they also endogenize weighted votes and study their role in an indirect democracy. The efficient voting rule is a mixture of weights and supermajority (Barberà and Jackson, 2006). Aghion, Alesina and Trebbi (2004) analyze the constitutional choice about the level of insulation of political leaders. The optimal degree of insulation depends on the cost of compensating the losers, the social benefits of policy reforms, the uncertainty about gains and losses, and the degree of risk aversion.

We differentiate our model from the above papers in many respects. First, we develop a model of constitutional design appropriate for the study of behavioral distortion in decision making. Second, we consider a situation where assembly members can have different subjective expectations, while in these models agents share the same ex-ante degree of uncertainty about policy outcomes. Third, we consider asymmetries in voting power, while these models generally rely on the "one head - one vote" assumption. Finally, since we use generic utility functions, we perhaps provide a more general treatment of risk aversion.

Some of the above-mentioned models describe the constitutional stage as a voting game (e.g. Messner and Polborn, 2004). Some conceptual questions, however, remain eluded in these models. For instance, when and why have agents agreed to vote on a constitution? How have they determined the threshold for voting on the voting threshold? We think unanimous bargaining is more appropriate to describe constitutional negotiations than formal voting on the voting rules. At the outset of any non-coercive constitutional process there must be a *unanimous* "voluntary exchange" in which everyone accepts the goal to establish a way to make common decisions.[3] We present the constitutional stage as a Nash unanimous bargaining game amongst *heterogeneous* members of a constituent assembly. We show that members who feel more subject to expropriation risk in future decisions have consistent leverage at the constitutional stage. This result links our model to Coelho's (2005). He obtains similar results in a setting where the voting rule is selected by the *normative* Rawlsian maximin criterion. He claims that this criterion is chosen on the basis of fairness. It is not clear, however, why this should be the case if voters are not behind a veil of ignorance. Why should the most conservative members accept a criterion that gives all the power to the least conservative voter? This question remains unanswered by some of the above mentioned models in which voting rules are selected by majority rule. We think that when voters are aware, at least partially, of their preferences, the axiomatic nature of the Nash bargaining approach makes it a superior modeling choice.

Finally, our work contributes to the recent and growing literature on behavioral political economy. We already mentioned the papers by Alesina and Passarelli (2015) and Ortoleva and Snowberg (2015). Bendor et al. (2011) present political models with boundedly rational voters/parties. Krusell et al. (2010) study government policies for agents who are affected by self-control problems. Lizzeri and Yariv (2015) study majority voting when voters are heterogeneous in their degree of self-control. DellaVigna et al. (2014) present, and experimentally test, a model of voter turnout with positive returns of voting on citizens' social image. Passarelli and Tabellini (2016) study how emotional unrest affects policy outcomes. None of these works addresses any constitutional issues, as we do in this paper. Recently, Bisin et al. (2015) presented a model of fiscal irresponsibility and public debt accumulation. Constitutional balanced budget rules, they claim, should restrain governments' response to voters' self-control problems. In their paper, the demand for more restrictive constitutional rules is related to time-inconsistency due to self-control. In our paper, the demand for higher super-majority is due to loss aversion.

The rest of the paper is structured as follows. Section 2 presents the setup of the legislative stage. It computes the optimal majority threshold and analyzes how it depends on individual features. Subsections 2.2 and 2.3 present the effects of loss aversion and overconfidence, respectively. Section 3 presents a constitutional stage in which we compute the equilibrium

[3] This is consistent with the "classical" approaches to constitutions of Wicksell, Lindahl, Musgrave, and many others. See Mueller (1973), for an excellent discussion.

threshold. Section 4 concludes. An Online Appendix contains all of the proofs of Propositions and Lemmas (Appendix A), and provides a discrete version of the legislative model, which is more suitable to describe a committee with a small number of voters (Appendix B).

2 The Legislative Lottery

Consider a set $N = \{1, ..., n\}$ of agents who make common decisions by voting. Let q be the majority threshold. w_i is agent i's number of votes ($i = 1, .., n$), and $m = \sum_N w_i$. Assume that $q > \frac{m}{2}$. The assembly has to deliberate on two exogenous policy proposals, α and β, and assume that N is partitioned in two subsets, the α-types and the β-types. If the adopted policy is α, then all the α-types gain with respect to the status quo, and the β-types lose. If β is adopted, the opposite is true. If no policy passes, then the status quo, ς, remains. Abstention is not possible: all the α-types vote in favor of α and all the β-types vote against it, and vice versa.[4]

With a slight variation in the meaning of variables, this "legislative" framework applies to electoral competition as well. In this case, α and β represent the *exogenous* electoral platforms proposed by two candidates to lead the executive branch of a government. Say an α-type citizen has incentive to reduce the ability of the β-candidate to pass platform β. This can be done by requiring a high super-majority to receive parliamentary approval. In Aghion, Alesina and Trebbi's (2004) terminology, this would mean that any future leader, whether β- or α-candidate, will be less "insulated".

Consider an agent j. Suppose she is an α-type, and let $u_j : \{\alpha, \varsigma, \beta\} \to \mathbb{R}$ be her utility function:[5]

$$u_j(\alpha) > u_j(\varsigma) > u_j(\beta) \tag{1}$$

We assume that agent types are private information, and p is j's subjective probability that any other agent in $N \setminus j$ is of type α and thus will vote for α.[6] Conversely, $(1 - p)$ is the probability

[4]Probably the simplest way to look at α and β is by considering them purely redistributive policies: α is a tax levied on β-types and totally transeferred to α-types, and vice versa. However, many other types of policy alternatives containing social or ideological aspects can be described in the same way.

[5]Notice that if $u_j(\alpha) \geq u_j(\beta) \geq u_j(\varsigma)$ the problem of finding the optimal threshold becomes trivial: agent j always (at least weakly) prefers the bare majority rule.

[6]In order to save on notation, p has not been indexed by j.

that any other agent will vote for β. In a way, p parametrizes how an agent feels similar/different to the others. For example, members of a small ethnic minority or an ideologically extreme faction are likely to have a low p. This parameter may also reflect an individual's psychological traits, and we will sometimes refer to it as j's degree of optimism. Instead of thinking of p as an exogenous parameter, one can alternatively interpret it as a Bayesian updating of priors, based on idiosyncratic signals (cf. Subsection 2.3).

Call $S_\alpha \subseteq N \setminus j$ the coalition of "other agents" who vote for policy α. Agent j's probability of winning is given by the probability that S_α collects at least $q - w_j$ votes. Agent j then "adds" her own w_j votes, and the majority forms. Given j's uncertainty, the sum of votes in S_α is a random event that behaves as the sum of $n-1$ independent random variables, Z_i, $(i = 1, ..., n;$ $i \neq j)$, where $Z_i = w_i$ with probability p, and $Z_i = 0$ with probability $(1-p)$. If the number of agents is sufficiently large the Central Limit Theorem applies.[7] Thus, the sum of votes is normally distributed with parameters $\mu_\alpha = \sum_{i \in N \setminus j} w_i p$, and $\sigma_\alpha^2 = \sum_{i \in N \setminus j} w_i^2 p(1-p)$. Let $f^\alpha(\cdot)$ be its density function. Similarly, the sum of votes for β behaves normally with parameters: $\mu_\beta = \sum_{N \setminus j} w_i(1-p)$ and $\sigma_\beta^2 = \sigma_\alpha^2 = \sigma^2 = \sum_{N \setminus j} w_i^2 p(1-p)$, whose density function is $f^\beta(\cdot)$. Then, j's subjective probability of winning, $\Pr_j \{\alpha, q\}$, is given by the probability that the sum of "favorable" votes lies in $[q - w_j, m - w_j]$. Her subjective probability of losing, $\Pr_j \{\beta, q\}$, is the probability that the sum of "unfavorable" votes lies in $[q, m - w_j]$:

$$\Pr_j \{\alpha, q\} = \int_{q-w_j}^{m-w_j} f^\alpha(x)dx \ , \qquad \Pr_j \{\beta, q\} = \int_{q}^{m-w_j} f^\beta(x)dx \tag{2}$$

Finally, j's subjective probability of maintaining the status quo is $\Pr_j \{\varsigma, q\} = 1 - \Pr_j \{\alpha, q\} - \Pr_j \{\beta, q\}$.

j's voting prospect can be described as a lottery, $L_j(q) = (\alpha, \Pr_j \{\alpha, q\}; \beta, \Pr_j \{\beta, q\}; \varsigma, \Pr_j \{\varsigma, q\})$, with three possible outcomes, $\{\alpha, \varsigma, \beta\}$, and attached subjective probabilities $\Pr_j \{\alpha, q\}$, $\Pr_j \{\beta, q\}$ and $\Pr_j \{\varsigma, q\}$.[8] The expected utility of this voting lottery is

$$EU_j(L_j(q)) = \Pr_j \{\alpha, q\} \cdot u_j(\alpha) + \Pr_j \{\beta, q\} \cdot u_j(\beta) + \Pr_j \{\varsigma, q\} \cdot u_j(\varsigma) \tag{3}$$

[7] With only a few number of voters (say less than twenty) the approximation of the Central Limit Theorem becomes quite large. Thus, a discrete model with exact probability distributions is more appropriate. We present it in Online Appendix B. The main results go through. However, we lose the benefits of differential calculus.

[8] A natural interpretation of this lottery is that in the legislative stage alternative α is posed against alternative β. This might sound strange if one usually thinks of the legislative process as a pairwise competition between

2.1 Optimal threshold

The optimal threshold q_j^* maximizes expected utility in (3). At an interior optimum the FOC is satisfied at a stationary point, q_j^0:

$$f^{\alpha}(q_j^0 - w_j) \cdot [u_j(\alpha) - u_j(\varsigma)] = f^{\beta}(q_j^0) \cdot [u_j(\varsigma) - u_j(\beta)] \tag{4}$$

Agents balance the marginal reduction in the expected benefits of belonging to the majority (the LHS of (4)) with the marginal reduction in the expected loss of falling into the minority (the RHS). Since $f^{\alpha}(\cdot)$ and $f^{\beta}(\cdot)$ are two known normal densities, it is easy to see that the unique stationary point is:[9]

$$q_j^0 = \frac{m}{2} + \frac{\sigma^2 \ln RASQ_j}{w_j + \mu_\alpha - \mu_\beta} \tag{5}$$

where

$$RASQ_j = \frac{u_j(\varsigma) - u_j(\beta)}{u_j(\alpha) - u_j(\varsigma)} \tag{6}$$

$RASQ_j$ is the Relative Advantage of the Status Quo, namely the ratio between j's benefits of not being tyrannized by an unfavorable majority, $u_j(\varsigma) - u_j(\beta)$, and the benefits of being part of a favorable majority, $u_j(\alpha) - u_j(\varsigma)$.

We say that j is "confident" about winning when, for any q, the chance of winning is always higher than the chance of losing (i.e. $\Pr_j\{\alpha, q\} > \Pr_j\{\beta, q\}$). Since $f^{\alpha}(\cdot)$ and $f^{\beta}(\cdot)$ have the same variance, $\Pr_j\{\alpha, q\} > \Pr_j\{\beta, q\}$ if and only if the mean of the former density *plus* j's votes is strictly larger than the mean of the latter one.

Definition 1 *Agent j is confident if $\mu_\alpha + w_j > \mu_\beta$. She is non-confident if $\mu_\alpha + w_j < \mu_\beta$.*

Confidence is related to pessimism, but is different. For instance, j may have a pessimistic view about how others will vote (i.e. $p < 0.5$, which in turn yields $\mu_\alpha < \mu_\beta$), but nonetheless she is confident about winning if she has sufficient voting power (so that $\mu_\alpha + w_j > \mu_\beta$). The following lemma shows the relationship between confidence and the concavity of $EU_j(L_j(q))$.

the current status quo and any proposal. However, no substantial changes would occur in the voting prospect if one assumes that, in a first round, any of the two alternatives (say α) is posed against the status quo ς; the winning one becomes the new status quo. Then, in a second round, the other alternative, β, is posed against the new status quo.

[9]See Online Appendix A for details.

Lemma 1 *i) If agent j is confident, then EU_j is concave for any $q \in \left[\mu_\beta, \mu_\alpha + w_j\right]$; moreover, EU_j is concave for any $q \in [q_s, m]$ if j is sufficiently confident (i.e., if $\mu_\alpha + w_j - \mu_\beta$ is positive and large enough).*

ii) If agent j is non-confident, then EU_j is convex for any $q \in \left[\mu_\alpha + w_j, \mu_\beta\right]$; moreover, EU_j is convex for any $q \in [q_s, m]$ if j is sufficiently non-confident (i.e., if $\mu_\beta - \mu_\alpha - w_j$ is positive and large enough).

By (3),

$$EU_j'(L_j(q_j)) = \frac{e^{-\frac{(q_j-\mu_\beta)^2}{2\sigma^2}}}{\sigma\sqrt{2\pi}}\left(RASQ_j - e^{\frac{(2q_j-m)(w_j+\mu_\alpha-\mu_\beta)}{2\sigma^2}}\right)$$

The sign of $EU_j'(L_j(q_j))$ is the same as the sign of the bracketed term, which in turn is determined by the difference between $RASQ_j$ and a function of j's degree of confidence.[10] For instance, if it is positive (negative) for any $q_j \in [m/2, m]$ we will have a corner solution: the agent will prefer unanimity (simple majority). When $EU_j'(L_j(q_j))$ is zero, a stationary point q_j^0 occurs (cf. expression (5)). The following proposition characterizes the optimal choice of q:

Proposition 1 *i) If j is confident,*
i.1) she prefers simple majority if $RASQ_j \leq 1$;
i.2) she wants a supermajority or unanimity if $RASQ_j > 1$.
ii) If j is non-confident,
ii.1) she prefers simple majority if $RASQ_j < \frac{y_\alpha}{(1-y_\alpha)} < 0.5$,
where $y_\alpha \equiv \Pr_j\{\alpha, q|_{q=q^s}\}$ is the probability of winning under simple majority;
ii.2) she prefers unanimity in all other cases.

To see the intuition, consider part *i)* of Proposition 1. It relies on the fact that $EU_j(L_j(q))$ is concave at q_j^0 when the agent is confident (cf. Lemma A.1 in the Online Appendix A). Part *i.1)* of the proposition says that if $RASQ_j \leq 1$ then $q_j^0 < q^s$, where q^s is the simple majority threshold.[11] In this case j prefers q^s as, intuitively, she thinks that, for any threshold, winning is more likely than losing. She then wants a rule that "eases" majority formation. Moreover, since $RASQ_j \leq 1$, the voting prospect yields an additional benefit: the cost of losing is (weakly)

[10] The second bracketed term, $e^{\frac{(2q_j-m)(w_j+\mu_\alpha-\mu_\beta)}{2\sigma^2}}$ is a monotonic tranformation of confidence, $w_j + \mu_\alpha - \mu_\beta$.

[11] $q^s = \left\lceil \frac{m}{2} \right\rceil$ if m is odd (where $\left\lceil \frac{m}{2} \right\rceil$ represents the rounding of $\frac{m}{2}$ up to the integer) and $q^s = \frac{m}{2} + 1$ if m is even.

lower than the benefit of winning. In this case she prefers the lowest possible threshold: the bare majority. Part *i.2)* says that the solution is interior (i.e. a supermajority) only if $RASQ_j > 1$ and $q_j^0 < m$. In this case a trade-off occurs: winning is likely, but losing is relatively costly, creating demand for protection. Thus a supermajority, if not unanimity, is preferred.

Take part *ii)* of Proposition 1. A non-confident voter becomes a risk-seeker. Her optimization problem always yields one of the two most risky values of q: namely, either simple majority or unanimity. Since losing is more likely than winning, she chooses simple majority if the relative advantage of winning is "sufficiently large". Specifically, if j is non-confident, $RASQ_j$ must be lower than the ratio of winning to losing probabilities under simple majority $y_\alpha / (1 - y_\alpha)$. Since this ratio is at most 0.5, then the condition in statement *ii.1)* is quite restrictive: a non-confident agent is relatively unlikely to prefer simple majority. In all other cases she chooses unanimity.[12]

2.1.1 Risk aversion, voting power and optimism

In general, an individual's preference for a majority threshold reflects the following features. First, a more risk averse agent prefers a higher threshold because majority formation can be blocked more easily. Her conservative attitude towards political changes translates into a stronger preference for less decisive voting rules. Second, voting power gives greater control over the collective decision, making the outcome more likely to be the preferred policy. An agent with more voting power is less conservative and wants to facilitate majority formation. Therefore, she prefers less protective voting rules. Third, optimism parametrizes an agent's subjective perception of being part of a majority (high p) or a minority (low p). As p decreases, demand for protection increases. Hence, the preferred threshold is increasing in p. The following proposition summarizes these results.

Proposition 2 *Agent j's most preferred threshold is*

i) (weakly) positively related to her degree of risk aversion;

[12] For completeness, Online Appendix A proves that if $\mu_\alpha + w_j = \mu_\beta$, then $EU_j(L_j(q))$ is linear. Therefore,

a) if $RASQ_j > 1$, then $EU_j(L_j(q))$ is increasing in q, hence j prefers unanimity;

b) if $RASQ_j < 1$, then $EU_j(L_j(q))$ is decreasing in q, hence j prefers simple majority;

c) if $RASQ_j = 1$, then $EU_j(L_j(q))$ is independent of q, hence j is indifferent: any q yields the same expected utility.

ii) (weakly) negatively related to her voting weight, w_j;

iii) (weakly) negatively related to her degree of optimism, p.

2.2 Loss aversion

Loss averse individuals perceive outcomes as gains and losses, relative to the status quo, and "*...losses loom larger than gains*" (Kahneman and Tversky, 1979; p. 279). These individuals display an endowment effect, i.e. a strong attachment to the status quo. In our model, loss aversion leads them to prefer higher thresholds as a way to increase the chance of maintaining the status quo.

Let $EU_j(L_j(q) \mid \varsigma)$ be the reference-dependent expected utility of individual j under loss aversion. By (1-3),

$$EU_j(L_j(q) \mid \varsigma) = \Pr_j\{\alpha, q\} \cdot [u_j(\alpha) - u_j(\varsigma)] - (1+\lambda)\Pr_j\{\beta, q\} \cdot [u_j(\beta) - u_j(\varsigma)] \tag{7}$$

where $\lambda > 0$ is the parameter which captures loss aversion. The first term in (7) represents the indirect benefit of winning, relative to the status quo, while the second term is the indirect cost of losing. This formulation satisfies the *decomposability* property: individuals bracket benefits and costs separately (Tversky and Kahneman, 1991; Köszegi and Rabin, 2006). In reference-dependent models, the way one defines the reference point is obviously critical.[13] In the present paper, the status quo represents quite a natural reference point since individuals look at the majority threshold as an instrument to lower the chance of changing the status quo. Maximizing (7) yields the following stationary point

$$q_{j\lambda}^0 = \frac{m}{2} + \frac{\sigma^2 \ln[(1+\lambda)RASQ_j]}{w_j + \mu_\alpha - \mu_\beta} \tag{8}$$

where the subscript λ denotes loss aversion. Comparing $q_{j\lambda}^0$ with q_j^0 defined by (5) tells us that j always prefers more protection when she is loss averse. This "demand" for more protection is increasing in the loss aversion parameter, λ.

Proposition 3 *Compared to the case without loss aversion,*

i) if j is loss averse and confident,

[13] See the recent literature on endogenous or forward-looking reference points (Köszegi and Rabin, 2006, and DellaVigna, 2009, for an extensive survey).

i.1) she is less likely to prefer simple majority: she does it only if $(1+\lambda)RASQ_j \leq 1$*;*

i.2) she wants a supermajority or unanimity if $(1+\lambda)RASQ_j > 1$*;*

ii) if j *is loss averse and non-confident,*

ii.1) she is less likely to prefer simple majority: she does it only if $RASQ_j < \frac{y_\alpha}{(1-y_\alpha)(1+\lambda)}$*;*

ii.2) in all other cases she prefers unanimity;

iii) the preferred threshold is less sensitive to j*'s voting power.*

Because of loss aversion, winning becomes more attractive relative to losing, which is psychologically costly. Loss aversion has then the same effect as an increase of $RASQ_j$ in the model without loss aversion, leading to stronger preference for protection. This result is consistent with the status quo bias emphasized by Alesina and Passarelli (2015) in a model with simple majority and heterogeneous loss averse voters. In that model there is no uncertainty, and a change occurs only if, for the majority of voters, the utility of a policy reform is sufficiently larger than the utility of the status quo. In the present model, individuals do not choose the policy, but the voting rule to choose the policy. Nonetheless, the effect of loss aversion reflects the same bias towards the status quo: a high supermajority implies that a constituency for a reform will exist only when a sufficiently large number of individuals will gain from that reform. In all other cases, the status quo remains.

2.3 Bayesian updating and overconfidence

We begin by considering how a rational agent updates her priors when she receives information about each other player's voting preferences. Then we study what happens if the agent is overconfident.

Bayesian updating

The variance σ^2 of the prior distributions $f^\alpha(\cdot)$ and $f^\beta(\cdot)$ measures how uncertain j's prior is. Suppose that j receives new information in the form of "how many votes have been cast for α and how many for β" in a number s_i of trials that involved a sufficiently large number of votes by voter i, for $i \in N \setminus j$.

This new information is then a sample of $s = \sum_{i \in N \setminus j} s_i$ draws from the distribution in which each voter makes her choice according to her true probability to vote for α and for β. We assume that the number of trials is large enough and it is the same for each voter, $s_i = s/(n-1)$

for any i. Let $s_{i\alpha}$ be the number of draws for voter i with outcome α, and $\bar{\mu}_k$ and $\bar{\sigma}_k^2 = \bar{\sigma}^2$ be the mean and the variance of sample observation regarding votes for k ($k = \alpha, \beta$) over s draws. Thus,

$$\bar{\mu}_\alpha = \sum_{i \in N \setminus j} w_i \frac{s_{i\alpha}}{s_i}, \quad \bar{\mu}_\beta = \sum_{i \in N \setminus j} w_i \frac{s_i - s_{i\alpha}}{s_i} \quad \text{and} \quad \bar{\sigma}^2 = \sum_{i \in N \setminus j} w_i^2 \frac{s_{i\alpha}(s_i - s_{i\alpha})}{s_i^2} \tag{9}$$

Both the sample means and the sample variance positively depend on voters' weights, w_i. The idea is that a given number of draws, say with outcome α, has a larger impact on the means and the variance if it regards a more powerful voter.

Since the distribution from which sample information has been drawn is normal,[14] Bayesian updating implies that the *posterior distributions*, call them $f^{\alpha|s}(\cdot)$ and $f^{\beta|s}(\cdot)$, are two *normals* as well, with the following parameters (Winkler, 2003):

$$\mu_{\alpha|s} = \frac{(\bar{\sigma}^2/s)\mu_\alpha + \sigma^2 \bar{\mu}_\alpha}{(\bar{\sigma}^2/s) + \sigma^2}, \quad \mu_{\beta|s} = \frac{(\bar{\sigma}^2/s)\mu_\beta + \sigma^2 \bar{\mu}_\beta}{(\bar{\sigma}^2/s) + \sigma^2} \quad \text{and} \quad \sigma_s^2 = \left(\frac{1}{\sigma^2} + \frac{s}{\bar{\sigma}^2}\right)^{-1} \tag{10}$$

These parameters imply that "bad news" (i.e., $\bar{\mu}_\alpha < \mu_\alpha$) yields two effects. First, j lowers her winning expectations: $\mu_{\alpha|s} < \mu_\alpha$. Second, she raises her losing expectations: $\mu_{\beta|s} > \mu_\beta$. "Good news" (i.e., $\bar{\mu}_\alpha > \mu_\alpha$) yields opposite effects. This downward/upward expectation revision is magnified when priors are: rather imprecise (high σ^2) compared to the precision of new information (high s and low $\bar{\sigma}^2$). As pointed out earlier, $\bar{\mu}_\alpha$ is more "reactive" to sample information $s_{i\alpha}$ coming from more powerful agents. Broadly speaking, this means that "bad news" ("good news") concerning the voting preferences of powerful agents are worse (better) than those concerning weak agents.

By (5) and (10) the stationary point after the signal is

$$q_j^{0|s} = \frac{m}{2} + \frac{\sigma_s^2 \ln RASQ_j}{w_j + \mu_{\alpha|s} - \mu_{\beta|s}} \tag{11}$$

Which voting rule will j prefer after the signal? Consider "good news" – the average sum of votes for α in the signal is larger than the prior average. In this case, j thinks that winning is more likely than she previously thought. The impact of the signal is the same as an exogenous increase in j's degree of optimism, p. Proposition 2 applies. Good news leads j to prefer a

[14] This distribution is Poisson binomial (sum of independent Bernoulli trials that are not necessarily identically distributed, due to different $s_{i\alpha}$). Since for each i the number of draws s_i is sufficiently large, the Central Limit Theorem applies. Thus, also a Poisson binomial can be approximated by a normal distribution (e.g., see Neammanee, 2005).

lower threshold. Larger variance of the priors and higher precision of the signal yield a larger impact of the signal on the most preferred rule. By Proposition 1, if j is confident she prefers a lower super-majority after the signal. In the case that she is non-confident, she eventually shifts from unanimity to simple majority.

The opposite occurs in the case of bad news: now j wants a higher threshold. The effect is the same as a decrease in optimism. The effect is strong when the negative signal is relatively precise compared to the prior. However, (11) shows that an interesting trade-off comes about when news is bad and the signal is very precise. On the one hand, an agent wants more protection because news is bad. On the other hand, she wants more decisiveness because, given the high quality of the signal, her level of uncertainty is lower after the signal. Inserting (10) into (5) tells us how she solves this trade-off.

Overconfidence

Overconfidence is a psychologically distorted reaction to new information. Existing literature has defined it in three different ways: *overestimation*, *overplacement*, and *overprecision* (cf. Moore and Healy, 2008). We consider overprecision because it is empirically robust and more general than overestimation or overplacement. Overprecision is an agent's attitude to think that the signal is more accurate than it actually is (e.g. Soll and Klayman, 2004, and Ortoleva and Snowberg, 2015). We model overprecision as an additional weight, $\phi > 0$, assigned to the number s of trials in the signal. Suppose an agent receives a signal consisting of s trials. Overprecision leads her to behave as if it consisted in $(1+\phi)s$ trials. Her posteriors are normally distributed with the following parameters, where the superscript "o" stands for overprecision:

$$\mu^o_{\alpha|s} = \frac{\frac{\bar{\sigma}^2}{(1+\phi)s}\mu_\alpha + \sigma^2\bar{\mu}_\alpha}{\frac{\bar{\sigma}^2}{(1+\phi)s} + \sigma^2}, \quad \mu^o_{\beta|s} = \frac{\frac{\bar{\sigma}^2}{(1+\phi)s}\mu_\beta + \sigma^2\bar{\mu}_\beta}{\frac{\bar{\sigma}^2}{(1+\phi)s} + \sigma^2}, \text{ and } \sigma^{o2}_s = \left(\frac{1}{\sigma^2} + \frac{(1+\phi)s}{\bar{\sigma}^2}\right)^{-1} \tag{12}$$

By (10) and (12), it is immediately apparent that $|\mu^o_{\alpha|s} - \bar{\mu}_\alpha| < |\mu_{\alpha|s} - \bar{\mu}_\alpha|$, and $\sigma^{o2}_s < \sigma^2_s$: the posterior of an overconfident agent is influenced too much by new information and has too small a variance. This leads the agent to be overly optimistic when she receives good news and overly pessimistic if news is bad. We pointed out earlier that if news is bad there is a trade-off between the sign and the quality of information. Online Appendix A proves that this trade-off disappears when the agent is sufficiently overconfident.[15] The impact of news on expectations is always larger than their impact on uncertainty. She always wants more protection when the

[15] By "sufficiently overconfident" we mean $\phi > \underline{\phi}$, where $\underline{\phi}$ is defined in the proof of Proposition 4.

news is bad, in spite of less uncertainty. In a sense, for an overconfident agent, the quality of information is less important than information itself. The following proposition summarizes these results.

Proposition 4 *Compared to the case with no overconfidence, if j is sufficiently overconfident,*
i) she is more likely to prefer simple majority or a lower supermajority if the signal contains good news about how the others will vote;
ii) she is more likely to prefer a higher supermajority or unanimity if the signal contains bad news.
iii) These over-reactions to information increase in the overprecision parameter ϕ.

The first statement is consistent with Ortoleva and Snowberg (2015). They show that overprecision leads to extremeness in political behavior. In their model, individuals want more radical reforms when they receive signals that lead them to think that more people share their same political preferences. This effect is fostered by overprecision. Our model "translates" the desire for more radical reforms into a desire for more decisive rules.

3 The Constitutional Game

Agents agree that voting can solve future conflicts between majorities and minorities. Making an agreement today about the method of making future decisions is more efficient than bargaining on every single future decision. This is consistent with reality and with a common approach to constitutions as incomplete contracts (Persson and Tabellini, 2000; Aghion and Bolton, 2002). We model the constitutional stage as a Nash bargaining game over a "material" outcome, which in this case is the majority threshold, q. There are at least two appealing features of this modeling choice. First, unanimity exposes negotiators to an implicit trade-off. On the one hand, it enhances the decisiveness of each negotiator: since no decision can be taken at the expenses of weak minorities, any valid proposal must adequately represent the interests of all negotiators. On the other hand, given the high costs of a failure, there is no incentive to adopt purely obstructionist strategies. Second, due to the neutrality and reasonability of its axioms, Nash bargaining can be adopted as a fair arbitration scheme that satisfies basic criteria of impartiality in distributive justice (Mariotti, 1999).[16]

As mentioned in the Introduction, some papers describe the constitutional stage as a voting game. For instance, if the constitutional stage adopted the simple majority rule, the equilibrium would be ensured only under the condition that all players are confident (because voters' preferences are single-peaked in this case). The equilibrium would be the median voter's most preferred threshold. Thus, all the power would rest on that voter. This huge concentration of power appears to be quite an unrealistic description of the constitutional negotiations.[17] One of the effects of the Nash bargaining solution is assigning some weight to all voters' preferences. This is why we referred to it as a more equitable and, perhaps, more realistic solution.

A standard assumption in constitutional analysis is that individuals are behind a veil of ignorance: they are unaware of any differences amongst each other. If this is the case, the issue of constitutional negotiations is empty: everyone agrees on the same voting rule. A non-trivial analysis of constitutional negotiations implies a certain degree of heterogeneity amongst agents.[18] In our perspective, heterogeneity may arise from risk aversion, degree of optimism, and voting power. While risk aversion is a subjective attitude, degree of optimism or voting power may reflect objective and stable differences in the constituents' original positions (e.g. ethnic minorities, poor regions in a federal country, small groups in a corporation, ...).

The choice of q at the constitutional stage will determine the voting lottery of the legislative stage. The payoff vector in the constitutional bargaining is the profile of the agents' expected utilities attached to the lotteries generated by q: $\{EU_1(L_1(q)), .., EU_n(L_n(q))\}$, where $EU_j(L_j(q))$ is defined by (3); $j = 1, .., n$.[19]

[16] A Supplementary Material available from the authors extends the cooperative constitutional bargaining model of this section. It includes a non-cooperative game of sequential bargaining *à la* Rubinstein in which n voters bargain over the majority threshold. It shows that when individuals tend to be patient, the solution of this game coincides with the Nash Bargaining Solution.

[17] See also footnote 22.

[18] "Constitutions are not written by social planners, and veils of ignorance have holes in them." (Aghion, Alesina and Trebbi, 2004, p. 578).

[19] Typically, constitutions design voting rules for "many" future legislative decisions. Thus, our previous analysis of the legislative lottery applies if we consider that α and β are not alternative proposals regarding a specific issue, but rather alternative future platforms or reforms in several different fields of the public life. Since at the constitutional stage there might be limited knowledge about the future, we can simply assume that gains and losses are opposed and equally sized values, say $\alpha = 1$, $\beta = -1$, and $\varsigma = 0$. This implies that at the