NBER WORKING PAPER SERIES

MARKET MECHANISMS FOR POLICY DECISIONS: TOOLS FOR THE EUROPEAN UNION

Alessandra Casella

Working Paper 8027
http://www.nber.org/papers/w8027

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2000

Paper prepared for the European Economic Association meetings, Bolzano, Aug 31-Sept 3, 2000.
The views expressed in this paper are those of the author and not necessarily those of the National Bureau of Economic Research.

Market Mechanisms for Policy Decisions: Tools for the European Union
Alessandra Casella
NBER Working Paper No. 8027
December 2000
JEL No. F33, H41, P16

ABSTRACT

The thesis of this paper is that more transparent, rule-bound and subtle mechanisms for policy coordination will be needed to ensure the success of an enlarged European Union. A common policy is a public good with distributional implications. Economists have developed a large number of plausible market mechanisms for the efficient provision of public goods, and the European Union, with its limited number of members and relative ease of information is a promising ground for such schemes. An important open area of applied research is thus the tailoring of incentive schemes to the specific needs of the European Union and its policy choices. The paper discusses two possible examples: a system of tradable deficit permits to implement the fiscal constraints imposed by the Maastricht treaty; and a rule allowing country representatives to shift their own votes intertemporally when deliberations are taken by vote in periodic committee meetings.

Alessandra Casella
Department of Economics
Columbia University
New York, NY 10027
and NBER
ac186@columbia.edu
tel: (212) 854-2459
fax: (212) 854-8059

INTRODUCTION

The European Union exists to coordinate the policies of its member states. Devising a coherent, common policy for a number of independent and diverse countries is a difficult matter, and will become more difficult still when the Union proceeds with the planned enlargement of its membership. The question of how best to achieve this coordination will remain the underlying thread behind the more specific debates facing the European Union. The thesis of this paper is that inspiration for more innovative and efficient answers can come from market games, the mechanisms studied by economists to induce optimal provision of public goods.

The observation is not very surprising. A common policy is a public good: as in the traditional example of national defense, an economic policy decision is identical for all regions and all individuals subject to it - think of the common interest rate and money supply in a monetary union, of common external tariffs in a custom union, of the common ceiling on fiscal deficits in the Pact for Stability and Growth. At the same time, not all regions and individuals are affected by the policy in the same manner or benefit from it equally. Thus in general a common policy is a public good shared by economic agents with heterogenous preferences. How to choose and finance such a public good in the most efficient fashion is a classical problem in economics, and economists have developed a large number of plausible market mechanisms that aim at eliciting truthful preferences and implementing the correct public good supply. The challenge then is to identify among the existing proposals those that could be tailored to the selection of macro policy in the European Union and realistically implemented.[1]

The first part of this paper describes briefly one such possible application: the efficient implementation of the fiscal limits established by the Pact for Stability and Growth through a system of tradable deficit permits. Inspired by the markets for environmental permits, the proposal suggests exploiting the incentives created by the possibility of trade to achieve fiscal

---

[1]From the original Clarke and Groves schemes, a multitude of applications and variants have developed. Of particular simplicity (and interest in our context) are market mechanisms for public goods provision in environment of complete information (for example, Bagnoli and Lippman (1989), Moore and Repullo (1988), Moore (1992), Varian (1994)).

discipline at minimum cost. Given the basic idea, the exact design of the scheme must take into account the specific features of the problem on hand, but with enough care and imagination it seems possible to address the concerns that the system may raise.

Market-based mechanisms have three main advantages over administrative rules: they rely on the creation of correct incentives, so that individual gains are aligned with the common good; they are flexible, because they allow freedom of choice; but at the same time they are rule-bound and transparent because markets require predictability of the rules. Motivated by these principles, but without advocating a full-fledged market, the second part of the paper discusses the possibility of making votes storable, when deliberations are taken by vote in periodic committee meetings. Without suggesting interpersonal exchanges of votes, the paper proposes allowing committee members to shift their own votes intertemporally, towards those times or decisions where the marginal return of voting is higher. Intuitively, the intertemporal equalization of marginal returns should translate into higher expected welfare, a conclusion supported by preliminary results but whose robustness will need to be qualified.

TWO APPLICATIONS

a. A Market for Deficit Permits[2]

The Pact for Stability and Growth, finalized in June 1997 as part of the Amsterdam Treaty, was designed to guarantee the fiscal health of the Union members. It advocates balanced budgets in the longer term and specifies, for each country, a ceiling for deficit spending of 3 per cent of GDP; a violation of the ceiling will trigger warnings and eventually penalties (unless exceptional circumstances can be invoked). The current cyclical expansion in Europe, coupled with the disciplinary measures that most member countries were forced to undertake prior to the establishment of the Euro, has improved the economies' fiscal positions and eased the concerns that accompanied the signature of the Pact. But the design of the Stability Pact suffers from

---

[2]This section draws on Casella (1999), where a more detailed analysis can be found.

important weaknesses, and these will become apparent and urgent again when the business cycle reverts phase, when countries tackle the structural reforms required by their pension systems, or when the challenges of enlarging the Union become realities.

The Stability Pact provides a clear example of the two twin limitations of the Union's current approach to policy coordination. First, the Pact lacks flexibility: in principle, barring extremely severe recessions, the limit to fiscal deficit applies to all countries at all times, independently of a country's cyclical phase, its structural conditions and the state of the European economy. Second, recognizing that the costs of enforcing the agreement could be at times unacceptably severe, exceptions are allowed. But what constitutes a legitimate exception is left intentionally vague, so that the rule's intransigence is softened by the uncertainty of the enforcement. The ambiguity of the rules will be more and more difficult to defend as new countries enter the Union, countries that do not share the decades of joint negotiations that have shaped current political equilibria within the Union.

But the objectives of the Stability Pact and the flexibility that the system needs are not incompatible: it is possible to design a scheme that promotes fiscal discipline, recognizes the realistic need for variable deficits and is at the same time transparent and predictable. Borrowing from the experience of environmental markets, we could transform the Stability Pact into a system of tradable deficit permits.

The starting point of the Stability Pact must be the belief that markets alone are unable to impose sufficient discipline on the fiscal position of individual members of the European Monetary Union. There are different interpretations of the origin of this market failure, but all have in common the belief that high deficit spending by any one country has costs for the others: a high deficit in Italy affects negatively Germany, France and the other members of the Union. Thus deficits are a form of pollution, originating in one country's activity but having repercussions for all. We can then adapt to our purposes the insights of environmental regulation.

Traditional pollution control takes the form of a ceiling on each source's emissions, typically identical for all. But because different sources utilize different production technologies, have different access to cleaner inputs and different scope for capital investment in pollution reduction, the same pollution limit imposes widely diverging costs. In a market for pollution permits, the regulatory authority sets the overall pollution limit - the total stock of permits available on the market - and the initial distribution of permits. Pollution sources sell permits if their marginal cost of pollution reduction is below the permits' price, and buy them otherwise. Cuts in pollution are distributed in such a way that all sources equalize their marginal costs of pollution reduction, achieving the target decrease in pollution at minimum total cost. The US market for emissions of sulfur dioxide, a contributor to acid rain, is the largest experiment of this kind. Established in 1990 but first binding in 1995, it targets coal-fired electric utilities in all US states and has succeeded in reducing emissions by more than 50 per cent of their 1980 levels, at savings estimated at about one third of total costs of traditional regulations for equal levels of compliance, or $300 million a year.[3]

The parallel with our fiscal problem is immediate: the current provisions of the Stability Pact consist of uniform quantitative constraints on each country's deficit, and the observance of this limit would be associated with very different costs, depending on the country's structure, debt overhang and cyclical phase. The dispersion in costs is a sign of the scheme's inefficiency: a system of tradable deficit permits would allocate deficits where their value is higher, making it possible to implement the desired fiscal discipline much more efficiently.

Consider how the scheme works in its simplest realization. Each year each country is allocated a number of deficit permits, equivalent, for example, to 3 per cent of its GDP. In practice, these permits could simply be entries in special accounts. The permits are denominated in Euros and freely tradable. At the time final fiscal statistics are made public, each country must have in its account a sufficient number of permits to cover the year's deficit, and these permits are withdrawn from the system. If a country is found not in compliance, it faces a steep fee (a fixed, very high price for each missing permit) and must relinquish a corresponding number of permits from the following year's allocation. Permits can be banked, and countries can prepare for anticipated shocks and buffer unanticipated ones, but cannot be borrowed from future allocations, to prevent governments from putting off needed adjustments. In practice, deficits

[3]Ellerman et al. (1997 a, b)

can be offset by permits carrying a date contemporaneous or preceding the year of the deficit. On-going or frequent auctions would allow governments to buy or sell permits according to their expectations of current and future needs.

Given the general idea of a system of tradable permits, the exact design can vary to reflect the specific policy concerns that have inspired the call for fiscal discipline. For example, countries with different debt positions can be treated differently, mirroring the fear that deficits from economies with larger outstanding debts may be particularly destabilizing for the Union as a whole. One plausible scheme would require countries to hold permits that are a multiple of their deficit, where the multiplicative factor is the country's debt to GDP ratio. Thus a country whose debt to GDP ratio is twice the size of the ratio of a second country would be required to hold twice as many permits for the same deficit, or equivalently would face a double price per unit of deficit. Over time, this scheme would lead the EMU countries to debt positions that are progressively smaller and more similar. And it would do so at the lowest aggregate cost.

A system of deficit permits would have a number of advantages. First, it would allow fiscal policy to respond to economic conditions - individual countries could choose to incur larger deficits by purchasing permits on the market. A negative shock specific to an individual country, the type of idiosyncratic shock that would put the existing system under stress, could be overcome at low cost, since the market price of the permits would reflect low demand by the other Union members.[4]

Second, the system would allow countries to undertake necessary long term reforms, even in the presence of short term negative repercussions on fiscal balances. Moving from an unfunded pensions system to a partially funded one, for example, could be accompanied by new debt issues - as it should be. In addition, a system of deficit permits would give countries incentives to tackle such structural problems when times are good (and permits cheaper), rather then relying on temporary measures and postponing reckoning.

Third, fiscal discipline would be achieved through a system of clear and predictable rules. Stating once more that flexibility and transparency are particularly important when countries are

---

[4]At the aggregate level, supply could be adjusted through predetermined automatic formulas in case of Europe-wide recessions.

more diverse may be redundant. But it is a point that we will do well to remember in the next few years.

B. A System of Storable Votes

The suggestion this paper advances is not to borrow mechanically economists' schemes and market solutions, but to draw inspiration from these schemes. In theory it would be possible to apply appropriate market designs and bidding mechanisms to a large number of common policy decisions, from coordinating fiscal positions to reaching agreement on the European Central Bank discount rate, to reforming the CAP. Many such mechanisms have been shown to possess desirable properties and some are simple enough that real world applications could at least be considered. But all share a common weakness - they rely on profit maximization and function through the use of money. When applied to interactions among national governments, these features may be problematic: governments' goals cannot be reduced to monetary gains, and although in many instances they can be aligned with them - monetary gains mean more resources to be used as the government desires - in other cases the simplification is not satisfactory. In addition, even in those applications where governments' trade-offs can be represented in monetary terms, the reality of credit markets imperfections implies that in market-based policy games countries' power would be determined by economic wealth. This hardly seems the correct principle on which to build new policy rules for a larger European Union. In some specific applications, for example when constraining fiscal imbalances through deficit permits, the problem can be corrected in part through the initial distribution of valuable resources - here the allocation of permits. But the difficulty is fundamental: incentive schemes depend on the exchange of a currency that is valued outside the specific decision being considered. Money is the most natural candidate, but it is ill-suited to political applications. Is there any alternative? This section discusses one possibility.

An important limitation of majority voting is the mechanism's inability to account for voters' intensity of preferences. In any decision-making period, each voter casts his vote (or the number of votes at his disposal) so as to maximize the chance of his preferred option winning, but the voter cannot cast more votes when he feels more strongly about the outcome because the number of votes is fixed. The implication is that the chosen allocation can often be improved upon. As stated in a classical textbook: "The condition for the Pareto optimality of the supply of public goods requires information on the relative intensity of preferences.[..] Since this information is not directly gathered under majority rule, it is not surprising that the outcomes under majority rule may not satisfy the Pareto optimality condition" (Mueller, 1989, p.82). In the exchange of private goods, on the other hand, individuals' responses to prices convey the necessary information. Thus it is very natural for an economist to think of designing voting rules where votes would function as prices, or more precisely would be equivalent to resources spent to bring about a desired outcome. To function in this manner votes must have value outside the specific decision being debated, and voters must be able to choose how many of these valuable votes should be cast on a specific issue. The classical scheme considered by the literature is logrolling - the possibility of trading votes. Under many circumstances, though not always, logrolling may indeed increase efficiency, but the scheme has weaknesses: first, and foremost, the explicit trade of votes is illegal in all democratic countries, and although more informal promises to exchange "voting favors" within committees are commonly used, they retain a moral stigma. Institutionalizing logrolling is not a feasible political reform of the institutions of the European Union.[5] Second, the scheme will lead to efficient outcomes in general only if all bargains are possible, "prices" are flexible and promises can be enforced. In a world where these conditions are not met, and the full Coasian bargain cannot be concluded, the first best is typically not attainable.

There is a natural alternative to vote trading that also would allow voters to express the intensity of their preferences, appears much less controversial, and has not been studied in the

[5]One moral value underlying our society is that there are goods that should not be bought and sold - and votes are among them. As for most moral values, this too is challenged at times: "Six people have offered to sell their votes for president on the Internet this week, fetching as much as $10,100 before the online auctioneer eBay canceled the bidding. [...] EBay said it was cooperating with investigators from the Justice Department." New York Times, August 19, 2000, p.A16. It would be interesting to know how the perspective buyers were planning to realize the trade.

literature: granting individuals the right to shift votes not interpersonally, but intertemporally. In other words, creating a system of storable votes, where decisions each period are taken according to the majority of the votes cast and not necessarily the majority of the voters. Consider for example a committee, meeting regularly to take some common decision - for example the Council of the European Central Bank deciding each month whether or not to change the value of the discount rate. Instead of having one vote to cast at each meeting, each committee member could be allowed to save his vote for future use, thus inducing him not to vote when relatively indifferent (or sure enough that his preferred option will be chosen anyway) and increase his relative weight in future deliberations. A similar voting scheme ("cumulative voting") has been analyzed in static problems where voters must choose a subset of several candidates to an office, and can divide their votes among the candidates in any proportion they wish.[6] But to my knowledge the intuitive idea of allocating votes over time has not received attention. Of course, the static and dynamic problem are identical in the absence of uncertainty, but it is exactly the progressive resolution of uncertainty, both about the voter's own evolving preferences and his opponents' voting behavior, that lends special interest to the dynamic scheme.

Given the constraint on interpersonal trading, there is no theoretical presumption that a system of storable votes should necessarily lead to Pareto-optimal allocations.[7] But my goals are more modest. As a first step, we should investigate whether the system is superior to static plurality rule in simple but plausible situations. Second, and more ambitiously, we can then enquire under what circumstances the ranking of the two rules applies more generally. Finally, once votes acquire value through the possibility of intertemporal transfers, the decision mechanism that maps the votes cast into a specific policy choice need not be restricted to traditional vote counting. For example, we could design particular auctions, tailored to replicate the properties of the best and simplest mechanisms for private provision of public goods that economists have proposed - but where voters bid with votes, as opposed to money.

---

[6]See, among others, Sawyer and MacRae (1962), Brams (1975), Guinier (1994), Gerber, Morton and Rietz (1998).

[7]Both logrolling and storable votes are 'second best' scheme. It would be interesting to compare their efficiency properties.

Intuitively, the possibility of shifting votes across time should allow individuals to smooth their utility from voting, or in other words to equate the expected marginal return of casting one's vote. This intertemporal arbitrage leads us to conjecture possible efficiency gains. But the analysis of the scheme is not trivial, because the number of votes cast will depend not only on the intensity of the voter's preferences, today and forecasted into the future, but also on the probability of affecting the final decision, and hence on each player's estimate of the other players' votes, today and into the future. To illustrate the intuition in the simplest possible manner, consider the following example, with two players and two periods.

Individuals i and j are faced on two consecutive periods by the same binary decision. As mentioned above, the decision could involve for instance modifying a common discount rate or leaving it as is[8], but for the purposes of this example I will just call the two alternatives between which the agents must choose *Blue* and *Red*. Ex ante, either alternative will be preferred by any individual voter with probability ½ and with varying intensity σ, where σ is drawn from a distribution $F(\sigma)$ defined over the support $[0,1]$ and is identically and independently distributed both across individuals and across times. In other words, if the policy decision does not match his preference, i's utility in period t is zero, but if it is his preferred alternative, then i's utility equals $\sigma_{i\,t}$. Individuals are unsure about their future preferences, but each period's preferences are realized before a vote is called, although an individual cannot observe the preferences of his opponent. Utility in period *2* is discounted by a factor δ. The distribution F is common knowledge.

Each individual has a total of two votes, to be distributed as he wishes over the two periods in support of his preferred decision. The alternative that receives most votes wins; if the votes are equal, the tie-break rule is a coin toss. Notice that a voter can never gain by lying about his preference and that in period *2*, the last period, he will always cast all remaining votes. Thus the problem he needs to solve is how many votes to use in period *1*.[9]

---

[8]Assuming, plausibly, that the direction of possible change is predetermined before voting takes place, and, less plausibly, that so is the magnitude.

[9]In this example, voters are allowed to move votes both backward and forward in time. With longer horizons, voters who might have run down completely their stocks of votes would

As shown in the Appendix, the optimal strategy is simple (and unique): in period *1*, a voter should cast *0* votes if the intensity of his current preferences σ_1 is lower than a threshold α and cast *2* votes if it is above, where, intuitively, α equals $\delta \int_0^1 \sigma dF(\sigma)$ and is strictly between *0* and *1*.[10]

We can now compare the expected value of this game at the design stage, before any preferences are realized, with the expected value of the traditional game where each player has and uses one vote each period. In this latter game, each period is identical ex ante, and since the two individuals always cast the same number of votes, individual *i* expects to win with probability *3/4* (with probability ½ *j* has the same preference; with probability ½ he does not, but *i'*s choice still wins if *i* wins the coin toss, and ½× ½=*1/4*). Thus if we call *W* the expected value of the one-period, non-storable votes game, and W_1 the value of the corresponding two-period game, we can write immediately :

$$W_1 = W + \delta W, \text{ where } W = 3/4 \int_o^1 \sigma dF(\sigma) \tag{1}$$

On the other hand, given the equilibrium strategy described above, the expected value of the storable votes game EV_1 is given by:

$$EV_1 = \int_0^\alpha \sigma dF(\sigma)\left(\frac{2+F(\alpha)}{4}\right) + \int_\alpha^1 \sigma dF(\sigma)\left(\frac{3+F(\alpha)}{4}\right) + \delta W \tag{2}$$

The details are in the Appendix, but once again the intuition is simple. With both voters following the same strategy, the possible loss from having fewer votes in the final period is

---

cause delicate political problems, and it may be preferable to allow votes to be stored but not borrowed.

[10]Dividing one's votes is never optimal. But the result is sensitive to details of the model specification.

exactly compensated from the possible gain of having more, and the expected value of the game in period *2* equals the expected value of the one-period game with equal non-storable votes. In period *1,* on the other hand, each individual knows that he will cast *0* votes whenever his realized σ is smaller than α (and *2* votes otherwise), and he weighs the expected utilities in the two ranges by the corresponding probabilities of his preferred option being chosen. (For example, if $\sigma_{i1} < \alpha$, voter *i* still obtains his chosen option if *j* agrees with *i*'s preference (with probability ½), or if *j* disagrees, but casts *0* votes and the coin toss favors *i* (with probability $½ \times F(\alpha) \times ½$). And correspondingly if $\sigma_{i1} \geq \alpha$).

Comparing (1) and (2), it is clear that the expressions are equal for $\alpha = 0$ or $\alpha = 1$, and it is easy to show that the value of the storable votes game must be higher for all α strictly between *0* and *1*.[11] The reason is obvious from (2): the possibility of storing votes allows individuals to shift the probability of winning from low to high σ realizations, i.e. towards those realizations of preferences that are more intensely felt. And this is indeed the intuition we wanted to capture.

There can be no presumption at this stage that the result will hold in general. I have found that in the case of two voters the conclusion does extent to an arbitrary horizon of *T* periods (and a slightly richer model). But in the case of multiple voters the length of the horizon becomes important and I conjecture that a storable votes system may prove more efficient only if the game is played over a sufficiently long number of periods.[12]

---

[11]In the interval $\alpha \in [0,1]$, the left-hand side of (2) is concave in α and has a unique internal maximum.

[12]The reason is actually quite interesting. Consider the example of three players. At an arbitrary time $t < T$, the value of holding *n* votes is superior to the value of holding *n-k*, for any distribution of votes held by the other two players. However, this is not true in the last period *T*: it is easy to see, for example, that if *i*'s opponents hold equal votes, player *i*'s probability of seeing his preferred option selected is identical for any positive available votes. Thus the marginal value of a vote in the last period is too low, and this distorts voting behavior as the game approaches its end. See Casella (2000).

CONCLUSIONS

The two proposals discussed in these paper are not likely to be panaceas. Even under the most optimistic evaluations, they do not aim at achieving the first best - tradable deficit permits minimize the costs of satisfying an exogenously given fiscal constraint, and storable but not tradable votes explicitly ignore some opportunities for mutual gains. But they are the type of simple, feasible, additional tools that may facilitate policy coordination in the European Union. Devising more schemes of this sort and understanding under which circumstances they can be valuable in practice is, in my opinion, a worthy research project.

References

Bagnoli, Mark and Barton L. Lipman, 1989, “Provision of Public Goods: Fully Implementing the Core through Private Contributions”, *Review of Economic Studies*, 56, 583-601.

Brams, Steven, 1975, *Game Theory and Politics*, New York: Free Press.

Casella, Alessandra, 1999, “Tradable Deficit Permits. Efficient Implementation of the Stability Pact in the European Monetary Union”, *Economic Policy*, 29, 323-361.

Casella, Alessandra, 2000, “Towards a Theory of Storable Votes”, mimeo, June.

Ellerman, A. Denny, Richard Schmalensee, Paul L. Joskow, Juan Paulo Montero and Elizabeth M. Bailey, 1997a. "Emission Trading Under the U.S. Acid Rain Program: Evaluation of Compliance Costs and Allowance Market Performance", Special Report, MIT, CEEPR, October.

Ellerman, A. Denny, Paul L. Joskow and Richard Schmalensee, 1997b. "1996 Update on Compliance and Emissions Trading Under the U.S. Acid Rain Program", MIT, CEEPR, November.

Gerber, Elizabeth R., Morton, Rebecca, B. and Thomas A. Rietz, 1998, ”Minority Representation in Multimember Districts”, *American Political Science Review*, 92, 127-144.

Guinier, Lani, 1994, *The Tyranny of the Majority*, New York: Free Press.

Moore, John, 1992, "Implementation, Contracts, and Renegotiation in Environments with Complete Information", in J. J. Laffont (ed.), *Advances in economic theory: Sixth world congress*. Cambridge University Press: Cambridge, 182-281.

Moore, John and Rafael Repullo, 1988, "Subgame Perfect Implementation" *Econometrica*, 56, 1191-1220.

Mueller, Dennis C., 1989, *Public Choice II*, Cambridge University Press: Cambridge.

Sawyer, Jack and Duncan MacRae, 1962, “Game Theory and Cumulative Voting in Illinois: 1902-1954", *American Political Science Review*, 56, 936-946.

Varian, Hal, 1994, "A Solution to the Problem of Externalities When Agents Are Well-Informed", *American Economic Review*, 84, 1278-1293.

Appendix

Consider player *i*, whose realized preference shock is σ_{i1} in favor of the *Red* alternative. The higher the number of votes he casts this period, the higher the probability that *Red* will be chosen, but the smaller his influence next period. If *v* is the number of votes spent in period *1* and
p_{jv} the probability that *j* casts *v* votes, then *i*'s expected utility when spending zero votes is:

$$Eu^i{}_1(v_i = 0) = \sigma_{i1}(3/4p_{j0} + 1/2p_{j1} + 1/2p_{j2}) + \delta\left(p_{j0}EV^i{}_2(2,2) + p_{j1}EV^i{}_2(2,1) + p_{j2}EV^i{}_2(2,0)\right) \quad \text{(A1)}$$

where $EV^i{}_2$ is *i*'s expected value of the game in the next and final period, as function of the votes available to *i* and *j* respectively. If the two voters choose the same number of votes, the probability that *i*'s preferred option is decided upon is *3/4*; if *j* casts more votes, then his preferred alternative is always selected, and that equals *Red* with probability ½. As for next period, if *i* plays *0* today, he will then have *2* votes available, while *j*'s votes will depend on *j*'s current strategy. Keeping in mind that *i* similarly always wins when he is the one casting more votes, we immediately derive *i*'s expected utility when he spends *1* and *2* votes:

$$Eu^i{}_1(v_i = 1) = \sigma_{i1}(p_{j0} + 3/4p_{j1} + 1/2p_{j2}) + \delta\left(p_{j0}EV^i{}_2(1,2) + p_{j1}EV^i{}_2(1,1) + p_{j2}EV^i{}_2(1,0)\right) \quad \text{(A2)}$$

$$Eu^i{}_1(v_i = 2) = \sigma_{i1}(p_{j0} + p_{j1} + 3/4p_{j2}) + \delta\left(p_{j0}EV^i{}_2(0,2) + p_{j1}EV^i{}_2(0,1) + p_{j2}EV^i{}_2(0,0)\right) \quad \text{(A3)}$$

Solving period *2*'s expected values, we can obtain an explicit solution to the optimal strategy. Since both players will use all their votes, recalling the probabilities of one's preferred option winning in the different voting scenarios, we derive immediately:

$$EV^{i}{}_{2}(s,k) = \begin{cases} \int_0^1 \sigma dF(\sigma) & \text{if } s > k \\ 3/4\int_0^1 \sigma dF(\sigma) & \text{if } s = k \\ 1/2\int_0^1 \sigma dF(\sigma) & \text{if } s < k \end{cases} \tag{A4}$$

Substituting (A4) in (A1), (A2), (A3), we derive the unique equilibrium strategy described in the text. Recalling that $W = 3/4\int_0^1 \sigma dF(\sigma)$ and using (A1), (A3) and (A4), we can then write:

$$\begin{aligned} EV_1 = {} & \int_0^{\alpha} \sigma dF(\sigma)(1/2 + p_{j0}/4) + \int_{\alpha}^{1} \sigma dF(\sigma)(3/4 + p_{j0}/4) + \\ & + \delta W[F(\alpha)(p_{j0} + 4/3p_{j2}) + (1 - F(\alpha))(2/3p_{j0} + p_{j2})] \end{aligned} \tag{A5}$$

Player *j* follows an identical equilibrium strategy. Hence $p_{j0} = F(\alpha)$, and $p_{j2} = 1\text{-}F(\alpha)$. Substituting these values in (A5), we obtain equation (2) in the text.