Learning aversion and voting rules in collective decision making.

Philippos Louis*

Abstract

We set up a model of collective decision making where the group can learn before making a decision. We fully characterize agents' behavior. We show how learning aversion can arise within some members of the group and show how this depends on the voting rule employed. In such an environment the optimal voting is a super-majority rule because it effectively reduces learning aversion and therefore improves the quality of decision making.

1 Introduction

Collective decision making differs from individual decision making in several ways. Here is one of them: while an individual would usually prefer to have as many information as possible available before making a choice, the same is not necessarily true for the members of a group. The individual decision maker can choose to use new information in order to update her beliefs about the "best" possible choice or ignore it if for any reason she thinks that it is better to do so. But receiving it cannot harm her because it is she that makes the final decision. On the other hand, members of a group have to take in to account not only how information will affect their own beliefs and preference ranking over alternatives, but also how it will affect their fellows in the group and consequently, how the change in beliefs and preference rankings may affect the final outcome. If it changes it in a way that is not of their liking some, unlike the individual decision maker, may prefer the group staying without the new bit of information. We call this learning aversion.

Whether a group member will be learning averse or not will depend, among other things, on the way the group is aggregating preferences in order to make decisions. In general, the more

*This is an ongoing project that is part of my PhD thesis realized in the IDEA Doctorate program of the Universitat Autònoma de Barcelona. I would like to thank Joan de Martí for the continued support and guidance.I also thank Antonio Cabrales,Enriqueta Aragones, Jordi Masso, Andrea Mattozzi, Matias Iaryczower and Pedro dal Bo for useful comments. Financial support has been provided by the Spanish Ministry of Education through the FPU fellowship program and project SEJ2005-01481.

his preferences can be reflected on the final outcome, the less learning averse we would expect him to be. In this paper we setup a model of collective decision making in which preferences are aggregated through voting. Learning aversion arises endogenously and we show how it is affected by the voting rule and how this interaction affects the quality of decision making.

Agents in our model must decide whether to implement a risky policy or not. As mentioned, the decision is taken by voting. Before voting the may receive some public and costless information that can resolve some of the uncertainty. Because of the heterogeneity in the preferences, some agents will be risk averse. Whether or not information is received will depend in a stochastic manner on how many members of the group are learning averse.

Our first main result is that more conservative voting rules will reduce learning aversion. This happens because with a more conservative voting rule, such as a super-majority instead of simple majority, more conservative agents are reassured that an a priori risky policy will be applied only if there is strong evidence that favors it and therefore welcome more information. If the voting rule is not conservative enough even weak evidence in favor of the risky policy may lead to adoption. In this case conservative agents will prefer no information at all.

Of course, for a given level of information a more conservative voting rule makes it less likely for new policies to be adopted. But the level of information is most likely dependent on a group's members incentives towards learning. If that is so, given our first result one can expect that in some circumstances a more conservative voting rule may induce more learning by the group and therefore yield less mistakes such as not adopting good policies. This is our second result.

More information can of course be costly and this is another reason why individually rational agents may prefer not acquire such information. What we want to point out is that individual incentives against obtaining information in collective decisions do not have to be only cost driven. While cost considerations can be the only reason against obtaining information in some cases, our analysis suggests that learning aversion could drive group members to exaggerate or even induce information acquisition costs. A party concerned with the quality of the decision making process would have to resort to different ways of inducing the group to learn in each case. We believe that our model can serve as a basis to understand individual incentives for information acquisition on the part of the group and plan to extend the paper in that direction.

This paper contributes to the economic literature on collective decision making and voting rules. [4] Buchanan and Tullock (1962) made a seminal contribution arguing that simple majority rules may allow “bad” projects to be implemented since it does not take in to account the costs to the minority. On the other hand, unanimity assures that only pareto-efficient projects are implemented. More recent authors have build on these intuitions offering a variety of arguments for and against different types of voting rules (see [10]Glazer (1989),[11] Guttman (1998),[1] Aghion and Bolton (2003),[3] Dal Bo (2006) and others). A comprehensive review of this literature is given in [12] Messner and Polborn (2008).

This last paper is also the most similar to ours in terms of setup and results. In their model agents do not know what their preferences will be in the future. They can either implement a policy right away or wait until their future preferences are known and then decide whether they want to implement it or not. They find that a more conservative voting rule (a super-majority) may induce less conservative behavior in the group. Unlike our paper, agents do not know other agents’ preferences and thus vote strategically (condition their behavior on being pivotal). Furthermore, the choices voted on in the two periods of the model are different: in the first round it is wait or implement, in the second it is implement or not. Waiting, which is equivalent to obtaining more information in our model, is directly compared to implementation. In our paper we separate learning from policy choices and leave the alternatives in the policy choice voting unaltered. This makes the voting rule in our model directly comparable to the one in other papers considering zero-one policy choices.

A very recent and active literature, driven by the work of [2] Austen-Smith and Banks (1996) and [5][6] Feddersen and Pesendorfer (1996,1998) on strategic behavior in information aggregation through voting, has now turned it’s attention to information acquisition ([13] Persico 2004,[8] Gerardi and Yariv 2008, [9] Gershkov and Szentes 2007). In these papers attention is given to the interaction between voting rules (and other features such as group size) and information acquisition. The main difference of our work with this literature is that we do not treat the collective decision making process as an information aggregation process, since all information is public in our setup. This also rules out strategic voting.

[7] Fernandez and Rodrik (1991) make an argument with a similar flavor to the one made here, although in a different setup. In their model agents vote on a reform which benefits some and hurts others. They show that a reform that may be ex-post beneficial to the majority may

not be implemented because of the individual uncertainty. If a reform is implemented but proves to harm the majority it can be rejected. This leads to reform aversion or status quo bias as they call it. We on the other side consider learning aversion. Thus, although the arguments and results appear quite similar there are fundamental differences. To see that consider again a single decision maker. While a single decision maker may or may not be benefited by a reform, and hence accept it or reject it according to her tastes, she will never be hurt by costless learning.

The remainder of the paper is structured as follows. The next section gives an example that illustrates the main argument on the existence of learning aversion and it's dependence on the voting rule. Section 3 presents the formal model and presents our first result concerning the relationship between learning aversion and the voting rule. In section 4 we show how under some conditions a super-majority voting rule can increase the quality of decision making. Section 5 concludes. Detailed proofs and calculations are given in the appendix.

2 The argument by an example

Our main argument is that in some situations individual members of groups involved in decision making may be learning averse and this will also depend on the decision making rule. To illustrate this point we give the following example.

It is Wednesday evening and Pau (a Catalan teenager) and his parents have to decide on dinner. None of them is in the mood to cook, so they can either stay home and order pizza or go to their favorite restaurant. Two issues are relevant on this specific evening: First, the restaurant may be closed. Second, F.C Barcelona is playing an important match. Pau, being an big Barça fan does not want to miss the match which is shown on TV. There might be a TV in the restaurant showing it, but he can not be sure. If there is he would love watching the game enjoying the juicy stake they prepare. Pau's father also looks forward to the juicy stake, but is not in the mood of walking the 3 blocks to the restaurant if he can not be sure it is open. Pau's mother really wants to get some air so she would particularly enjoy the walk to the restaurant even if would turn out to be closed. The following table summarizes Pau's family payoffs[1] :

[1]The argument can be made without assigning specific payoffs, but they do help the illustration. In general, though, I would agree with the big credit card company: An evening stroll in Barcelona, watching a nice football match on TV and enjoying a juicy steak are all priceless.

	Staying home	Rest. closed	Rest. open, no TV	Rest. open, with TV
Pau	5	-100	-100	6
Father	0	-20	10	11
Mother	0	0	10	12

The uncertainty about whether the restaurant is open as well as whether it has a TV showing the match can be solved by making a phone-call. Let's suppose Pau's family always makes decisions that involve a high cost[2] unanimously. If they call the restaurant and find out the match is shown they will all agree to go. If the restaurant is open but the match not shown, Pau will object and the family will stay home. If it is closed they will agree to stay home. Hence Pau is strictly better off if the phone-call is made.

Now suppose the family makes the decision using simple majority. Also suppose that the odds for the restaurant being open are such that given no more information the father prefers staying home than risking a fruitless evening walk. Then, if they do not call the restaurant Pau and his father vote for staying home and that is what they do. If they call the restaurant and it is closed they will all agree to stay home. If it is open, both parents vote for going to the restaurant. In case the match is not shown, going to the restaurant is a disaster for Pau. He is therefore better off if nobody calls the restaurant to find out whether it is open. The best way for him to be sure to watch the match is to cut the phone line!

It should already be clear: using a voting rule that gives less power to the individual (simple majority) can make some individuals (Pau, in our case) learning averse.

3 The model

We now define a formal model of collective decision making. A continuum of agents is uniformly distributed in the policy space, which is the interval from 0 to 1. Hence, each agent is characterized by her type $h_i \in [0, 1]$. The state of nature is $\theta \in \{0, 1\}$. Both states of nature are a priori equally likely.

The group of agents must choose a policy from the set of alternatives: $A = \{0, 1\}$. Before

[2] Juicy stakes do not come cheap

choosing a policy the group may obtain further information about θ. This information, if obtained, takes the form of a public signal $s \in [0,1]$ drawn randomly from a distribution $f(s|\theta)$. We assume that $f(\cdot|\theta)$ is a triangular distribution over $[0,1]$ with the mode coinciding with θ. Following this formulation we have that $Pr(\theta = 1|s = \hat{s}) = \hat{s}$ and $Pr(\theta = 0|s = \hat{s}) = (1-\hat{s})$. The triangular distribution for the public signal is not necessary for our results to hold. Any other distribution for which $Pr(\theta = 1|\hat{s}) > Pr(\theta = 1|\hat{s}')$ when $\hat{s} > \hat{s}'$ and $Pr(\theta = 0|\hat{s}) > Pr(\theta = 0|\hat{s}')$ when $\hat{s} < \hat{s}'$, should yield similar results. The same is true for symmetry. Using the symmetric triangular distributions simplifies exposition. Specifically note that under this specification the ex-ante distribution of the public signal is uniform in $(0,1)$.

In order to simplify exposition we assume that there are two rounds and the group may decide on a policy either in the first round, $t = 1$, without obtaining a public signal or in the second round, $t = 2$ after obtaining the public signal. We will therefore refer to obtaining information as "waiting". Whether the group will wait or not is chosen randomly in a way that depends on agents' preferences over waiting. In particular we will assume for now that the group will wait with a probability μ_w which is equal to the mass of agents that prefer waiting. One interpretation for that is that agents that prefer waiting make some effort that increases the probability that the group will wait. The policy is finally chosen by voting following a quota rule. In particular, policy 1 is chosen over the status quo policy 0 if a proportion of at least $k \in [\frac{1}{2}, 1]$ votes for 1. In other words, if $k = 1$, policy 1 is chosen only under unanimity and $k = \frac{1}{2}$ corresponds to simple-majority.

Given the context of collective choice one might expect the group to simply vote whether to wait or not. The stochastic choice can be seen as similar to voting in a model where instead of a continuous mass of agents there is a discrete number of them chosen randomly. That specification would also result to the group waiting with a certain probability which would depend on the distribution of agents over the policy space. Still, the reason we choose this formulation goes beyond this equivalence.

Even if a group takes a vote to decide whether to seek more information on an issue or not, the final result, that is the quality of the information received often depends heavily on the actions

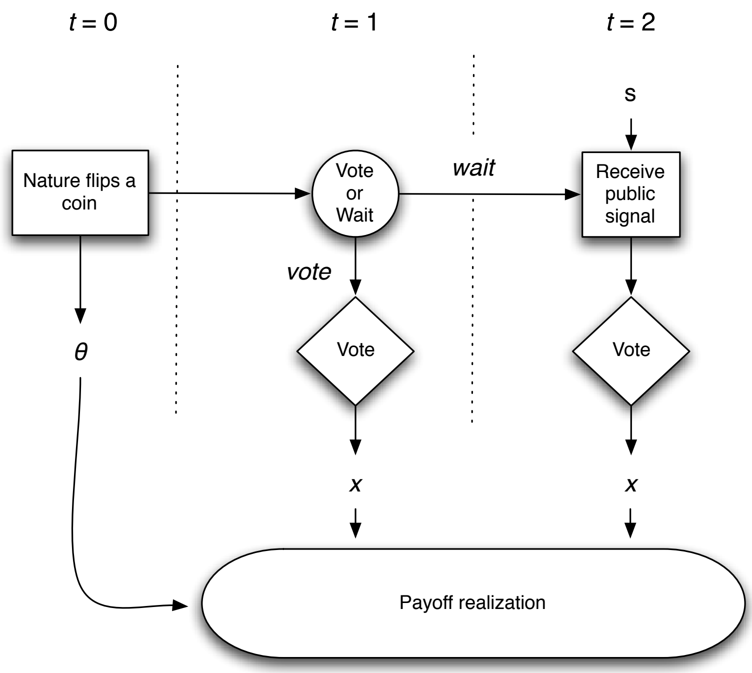

Figure 1: Timeline of the model

of the group members. Members of the group which are learning averse can make less effort to help information acquisition, commit less resources or even sabotage the process by creating confusion[3]. Hence, even if a vote is taken, the final outcome is not likely to be deterministic and will depend on the number of learning averse group members. Of course, our model does not allow for agents behaving in the way described, but that is one of the directions towards which we hope to expand the model in the future.

Agents care for two things. On the one hand they want a policy close to their own policy type. On the other hand they want the groups choice to match the true state θ. The agents' payoffs are given by

$$u_i(x) = -\alpha(h_i - x)^2 - (1 - \alpha)C(x, \theta)$$

[3]We have chosen the word waiting as a substitute for obtaining information. Still, waiting can be one of the methods applied by learning averse group members: they could suggest appointing a subcommittee with the task to gather the relevant information. At the same time they could make sure the subcommittee would not dispose of the necessary resources to perform the task. After "waiting" for a long time, the subcommittee will finally hand-in an inconclusive report. In the meantime, the new project may not be feasible or relevant any longer.

$$C(x,\theta) = \begin{cases} 0\,,\ x=\theta=1 \\ c\,,\ x=0 \\ 1\,,\ x=1\,,\ \theta=0 \end{cases} \quad ,\ 0 < c \leq \frac{1}{2}$$

where x stands for the choice made by the agents and the parameter $\alpha \in (0,1)$ indicates the degree of importance of the idiosyncratic part of an agent's preferences with respect to the "social" part.

This payoff structure is quite simplistic. Still it captures the common trade-offs between adopting new policies or remaining at the status quo. Using a more general form where payoffs from choosing the status quo are allowed to differ depending on the state of nature would yield similar results at the expense of expositional clarity.

Concerning individual types, there are different ways of interpreting this feature. The direct way would be to view it simply as bliss points in some kind of a policy space, say left-right. An alternative interpretation could be that individuals are conservative in different degrees. That is that they are not willing to endorse a new policy unless their is strong evidence in favor of it. How strong this evidence needs to be depends on how conservative they are. According to this interpretation, the closer an individuals type is to zero, the more conservative she is.

All information related to the game including other agents' preferences is common knowledge.

Let us propose an example of the "real-life" situations we wish to capture with this model. Consider a community council that must decide whether to approve the building of a new shopping mall (the new project: policy 1). The developers promise that it will create jobs and offer entertainment to all members of the community ($C(x=1,\theta=1)$). Still, there are strong concerns on the effects the new mall can have on existing local businesses as well as on the urban environment , since it would occupy one of the few open spaces left in the particular part of the community ($C(x=1,\theta=0)$). Therefore, some suggest the community is better-off without a new mall ($C(x=0,\theta)$). Some members of the local university have offered to prepare a study of the possible economic and environmental effects from building the shopping mall (s). The community would only need to put pressure on the dean or the central government to provide the necessary funds (μ_w).

Agents behavior is characterized by whether the agent prefers the group waiting and whether she votes for the status quo or for the new policy, conditional on whether the group waited or not: the waiting behavior and the voting behavior. An agent's waiting behavior will depend on all individuals voting behavior. Therefore, in order to determine the former one must first examine the latter.

3.1 Voting behavior

Once agents face the final vote their best response is to vote for the policy that if chosen by the group yields the highest expected payoffs. Unlike models of the latest Condorcet jury literature, agents here have no private information. They cannot infer any more information from other agents' voting behavior. Since their vote will only count if they are pivotal, they should then vote for their preferred (conditional on available information) policy. A vote may take place either in $t = 1$ or in $t = 2$. The difference is that in $t = 2$ agents receive the public signal before voting. Let us first look at the simpler case in $t = 1$.

Voting in $t = 1$

The payoff from choosing policy 0 for agent i is:

$$u_i(0) = -\alpha h_i^2 - (1-\alpha)c \tag{1}$$

The expected payoff from choosing policy 1 is:

$$E_{t=1}[u_i(1)] = -\alpha(1-h_i) - (1-\alpha)\frac{1}{2}$$

By comparing the two an agent can determine which policy yields a higher expected payoff. Following the previously explained rationale that is the policy the agent votes for. Obviously $u_i(0)$ is decreasing in h_i while $E_{t=1}[u_i(1)]$ is increasing. There must be an agent for which both policies yield the same expected payoff. That will be the indifferent voter for $t = 1$.:

The indifferent voter is $i^* = \{i : u_i(0) = E_{t=1}[u_i(1)]\}$:

$$h_{i^*} = \frac{1}{2} + \frac{1-\alpha}{2\alpha}\left(\frac{1}{2} - c\right)$$

All agents to the left of i^* vote for the s.q. 0. All agents to the right vote for policy 1. Note that at least half of the agents vote for the s.q. regardless of the values of α and c. Thus as long as $c < \frac{1}{2}$, the s.q is always chosen by the group if the vote takes place in $t = 1$ (remember that $k \geq \frac{1}{2}$).

Voting in $t = 2$ after observing $s = \hat{s}$.

Similar calculations are made in order to infer the agents' voting behavior if the group waits and votes in $t = 2$. The difference is that before the vote the public signal s is received. After receiving s agents update their beliefs about θ. Recall that $Pr(\theta = 1|s = \hat{s}) = 1 - Pr(\theta = 0|s = \hat{s}) = \hat{s}$. The payoff from choosing the s.q. is the same as in $t = 1$ and given in 1: The expected payoff from choosing policy 1 is:

$$E_{t=2}[u_i(1)|s = \hat{s}] = -\alpha(1 - h_i) - (1 - \alpha)(1 - \hat{s})$$

As before, $u_i(0)$ is decreasing in h_i while $E_{t=2}[u_i(1)|s = \hat{s}]$ is increasing.
The indifferent voter is $i^{**} = \{i : u_i(0) = E_{t=2}[u_i(1)]\}$:

$$h_{i^{**}} = \frac{\alpha c - c - \hat{s} + \hat{s}\alpha + 1}{2\alpha} =$$

$$h_{i^*} + \frac{(1 - \alpha)}{2\alpha}\left(\frac{1}{2} - \hat{s}\right)$$

Again, all agents to the left of i^{**} vote for the s.q. 0. All agents to the right vote for policy 1. If $\hat{s} = \frac{1}{2}$, then the indifferent agents in$t = 2$ is the same as in $t = 1$. This is natural, since $\hat{s} = \frac{1}{2}$ means that the public signal does not contain any new information. The higher $\hat{s}$, the lower $h_{i^{**}}$, or else the more agents that vote for policy in $t = 2$. This again has a natural explanation: the stronger the evidence in favor of policy 1, the more agents are willing to vote for it. How strong must the signal be for the group to choose policy 1? Policy 1 is chosen if:

$$1 - h_{i^{**}} < k$$

Solving this for the signal $\hat{s}$ gives us a threshold for the public signal. If $\hat{s}$ is above that threshold, policy 1 is chosen. If not, the status quo remains. The threshold is:

$$\tilde{s} = 1 - c + \frac{\alpha}{1-\alpha}(2k - 1) \tag{2}$$

The threshold $\tilde{s}$ is increasing in k. This is quite intuitive: a higher k means a more conservative voting rule. Stronger evidence is needed in order to convince the necessary amount of group members to vote the new policy. As long as the following condition is true, the threshold will exist.

$$\alpha < \frac{c}{c + 2k - 1} \tag{3}$$

3.2 Waiting behavior

Given the voting behavior of all agents as specified above, group members know the outcome of the vote in each round (conditional on the public signal for $t = 2$). Therefore agents calculate ex-ante expected payoffs of the two possibilities: the group waiting or not. From this calculation they compare these expected payoffs and determine whether they prefer waiting or not.

For agent i, the expected payoff of not waiting is equal to the payoff from choosing the s.q., given in 1, since this is the result if the group votes in $t = 1$. The expected payoff of waiting for i is:

$$E[u_i(\cdot); waiting] = \int_0^{\tilde{s}} u_i(0)ds + \int_{\tilde{s}}^1 E_{t=2}[u_i(1)|s]ds$$

Comparing the two allows us to calculate the indifferent "waiter":

$$i^w = \{i : u_i(0) = E[u_i(\cdot); waiting]$$

The indifferent "waiter's" type is:

$$h_{i^w} = \frac{3}{4} - \frac{k}{2} - \frac{1-\alpha}{\alpha}\frac{c}{4} \tag{4}$$

All agents to the left of h_{i^w} prefer not to wait. As mentioned, the probability that the group waits μ_w equals the mass of group members that prefer waiting. Thus:

$$\mu_w = 1 - h_{i^w} \tag{5}$$

Proposition 1. *A more conservative voting rule (higher k) increases the probability that the group will wait.*

Proof. Simply substitute 4 in to 5 and note that μ_w is linearly increasing in k. □

The intuition behind proposition 1, is quite clear. Take a conservative group member i, that is with h_i relatively close to 0. She is willing to vote the new policy if strong evidence appears in favor of it. If the voting rule k is low and the group waits, then it is possible that the signal is not strong enough to convince i, but enough other less conservative group members are convinced and policy 1 adopted. Agent i therefore prefers the group not to wait. If a higher k is set, then i can be sure that the policy will be adopted only if evidence in favor of policy 1 is strong enough. She will therefore prefer waiting. A higher k turn agent i from learning averse to non-learning averse.

4 Decision making quality

In this context, decision making quality refers to taking the "right" decision, that is the decision that matches the state of nature θ. The group should avoid making errors. Individual preferences could also matter, but in this setup the symmetric uniform distribution of types ensures that gains or losses from making one choice over the other will cancel out across individuals. There are two types of errors possible: i) adopting bad projects and ii) not adopting good projects. These bare different costs and are affected differently by the voting procedure. Suppose that the vote takes part in $t = 2$ with exogenous probability $\mu_w = 1$. Then, a more conservative voting rule makes it less likely for the new policy to be adopted. Errors of the first type are reduced, while errors of the second type increased. On the other hand, if the exogenous probability is $\mu_{w=0}$, the vote takes place in $t = 1$ and the s.q. is chosen. Errors of the first type never happen, while errors of the second type happen half of the time. Still, in our model the probability μ_w to vote in $t = 2$ is endogenous and as we showed in proposition 1, it is an increasing function of the voting rule. So, increasing k increases the probability of voting in $t = 2$ instead of $t = 1$ and this has a positive effect on type i) errors. On the other hand, the rule now is more conservative and this has a negative effect on type i) errors. The opposite effects are in place for type ii) errors.

We now look at what is the optimal voting rule given this trade-off. This depends also on the parameters α and c of the payoff function, for two reasons: First, different costs and gains are associated with different errors or correct decisions respectively and these depend on the parameter c. Second, the voting and waiting behavior of agents and thus the probability of making a good or bad decision depend on both parameters.

Definition 1. Let:

$$W(k; \alpha, c) = -(1 - \mu_w)c - \mu_w \left[\int_0^{\tilde{s}} c ds + \int_{\tilde{s}}^1 (1 - s) ds \right] \tag{6}$$

Represent the social welfare function.

The first term represents the payoff from not waiting and choosing the s.q. The second term represents the expected payoff from waiting. Inside the brackets the first term is the payoff from choosing the s.q. after receiving the public signal while the other one is the expected payoff from making an error of the first type. Remember that choosing a good new policy has zero payoff.

The optimal voting rule should maximize the social welfare function $W(k; \alpha, c)$

Proposition 2. *The optimal voting rule $k^*(\alpha, c)$ has the following properties:*

1. $k^*(\alpha, c) \in (\frac{1}{2}, 1]$
2. $\frac{\partial k^*(\alpha, c)}{\partial c} > 0$
3. $k^*(\alpha, c) = 1\,,\ for\ \alpha < \bar{\alpha}(c)\,,\ \frac{\partial \bar{\alpha}(c)}{\partial c} > 0$

Furthermore, it is the rule that would be chosen ex ante by the median type.

According to this proposition, the optimal voting rule is a super-majority rule. This is in part due to the higher expected cost of the new project. Notice that if learning was exogenous($\mu_w = 1$) and $c = \frac{1}{2}$, it is easy to see that the optimal voting rule is simple majority. Here instead, endogenous learning creates the effect of the voting rule on learning aversion. A more conservative voting rule induces less learning aversion which makes receiving the public signal more likely. Therefore, even in the case where the new project offers the same expected payoff as the status quo, a super-majority rule increases the quality of the decision making. If the rule becomes overly conservative, and the idiosyncratic part of the preferences is too strong, the group rejects projects that are likely to be good.

The last statement, concerning the median, comes from the fact that types are distributed uniformly across the whole policy space. It is still interesting though to notice that even in such a highly stylized model, if agents were allowed to choose the constitution choosing the condorcet winning voting rule, that would be a as super-majority.

5 Conclusions

This paper takes a novel perspective on collective decision making and the interaction between group learning and the voting rule. We explain the nature of this interaction and show how standard results in the literature on voting rules are affected once the possibility of learning aversion is taken in to account.

Having established that it would be interesting to examine in more detail the mechanisms through which learning aversion is manifested in groups. Or more generally, the way in which groups learn and how that learning is affected by individual actions of the members. With the general framework in place, this would consist in replacing the endogenous probability of learning with some kind of "information production" function. Group members would provide the inputs in such a function, and the possibility of sabotage could also be taken in to account. Although the flavor or our results would be preserved in such a setting, more concrete statements could be made regarding the behavior of group members in real life situations. We are working on such an extension.

References

[1] P. Aghion and P. Bolton. Incomplete social contracts. *Journal of the European Economic Association*, Jan 2003.

[2] D. Austen-Smith and J. Banks. Information aggregation, rationality, and the condorcet jury theorem. *The American Political Science Review*, Jan 1996.

[3] E. D. Bo. Committees with supermajority voting yield commitment with flexibility. *Journal of Public Economics*, Jan 2006.

[4] J. Buchanan and G. Tullock. The calculus of consent. *books.google.com*, Jan 1962.

[5] T. Feddersen and W. Pesendorfer. The swing voter's curse. *The American economic review*, Jan 1996.

[6] T. Feddersen and W. Pesendorfer. Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *The American Political Science Review*, Jan 1998.

[7] R. Fernandez and D. Rodrik. Resistance to reform: Status quo bias in the presence of individual- specific uncertainty. *The American economic review*, pages 1146–1155, Dec 1991.

[8] D. Gerardi and L. Yariv. Information acquisition in committees. *Games and Economic Behavior*, 62(2):436–459, Mar 2008.

[9] A. Gershkov and B. Szentes. Optimal voting schemes with costly information acquisition. *Journal of Economic Theory*, page 33, May 2008.

[10] A. Glazer. Politics and the choice of durability. *The American economic review*, Jan 1989.

[11] J. Guttman. Unanimity and majority rule: the calculus of consent reconsidered*. *European Journal of Political Economy*, Jan 1998.

[12] M. Messner and M. Polborn. The option to wait in collective decisions. 2008.

[13] N. Persico. Committee design with endogenous information. *Review of Economic Studies*, Jan 2004.

APPENDIX

Proof of Proposition 2 We are looking for the k that maximizes social welfare as defined in 6:

$$W(k;\alpha,c) = -(1-\mu_w)c - \mu_w\left[\int_0^{\tilde{s}} cds + \int_{\tilde{s}}^1 (1-s)ds\right]$$

By substituting μ_w and $\tilde{s}$ from 5 and 2 respectively and rearranging we obtain:

$$W(k;\alpha,c) = \frac{-c^3(\alpha-1)^3 + c^2(2k+1)\alpha(\alpha-1)^2 + c\alpha((4(k-1)k-7)\alpha+8)(\alpha-1) - (1-2k)^2(2k+1)\alpha^3}{8(\alpha-1)^2\alpha}$$

This is a third degree polynomial in k with a negative factor for k^3. Notice that:

$$\frac{\partial W}{\partial k}(k=\frac{1}{2}) = \frac{c^2}{4} > 0$$

Hence, the polynomial's local maximum obtains for some $k^* > \frac{1}{2}$. Also notice that:

$$\frac{\partial W}{\partial k}(k=1) = \frac{c^2(\alpha-1)^2 + 2c\alpha(\alpha-1) - 7\alpha^2}{4(\alpha-1)^2} < 0$$

$$\Rightarrow$$

$$\alpha > \bar{\alpha}(c) = \frac{c^2+c}{c^2+2c-7} + 2\sqrt{2}\sqrt{\frac{c^2}{(c^2+2c-7)^2}}$$

Thus, if α is lower than this threshold we have a corner solution, since $k=1$ is the maximum value for the voting rule. Finally, by taking the partial derivative and equating to zero we obtain the optimal voting rule:

$$k^*(\alpha,c) = \frac{c\alpha^2 + \alpha^2 - c\alpha + 2\sqrt{c^2\alpha^4 - c\alpha^4 + \alpha^4 - 2c^2\alpha^3 + c\alpha^3 + c^2\alpha^2}}{6\alpha^2}$$

The statement about the median type is derived by noting that given that $h_m = \frac{1}{2}$ the idiosyncratic payoff for this agent is the same for both policies. Therefore, his ex-ante expected payoff is a linear transformation of $W(k;\alpha,c)$ and therefore is maximized for the same value of k.