

Learning in Elections and Voter Turnout

Stefano Demichelis* and Amrita Dhillon[†]

July 2009

*Department of Mathematics, University of Pavia, email: sdm.golem@gmail.com

[†]Address for Correspondence: Amrita Dhillon, Department of Economics, University of Warwick, Coventry CV4 7AL, UK. Email: A.Dhillon@warwick.ac.uk. We would like to thank Alan Beggs, Thomas Palfrey, Matthew Jackson, V.Bhaskar, Myrna Wooders, Jonathan Cave, Michel Le Breton, Rabah Amir, Debraj Ray, and seminar participants at Ischia and the universities of Cambridge, Exeter, Essex, Warwick Mathematics, Montreal and Toulouse, Pennsylvania State University, TAM Mexico, New York University, Florida International University, for comments.

Learning in Elections and Voter Turnout

Abstract

Voter turnout in game theoretic models of voting has typically been difficult to predict because of the problem of multiple Nash equilibria. Many of these equilibria require an extreme precision of beliefs among voters that is unlikely to be reached in real elections. At the same time, mechanisms like pre-election polls exist to shape the beliefs of voters about expected turnout. We combine these two features in a model of voter learning in elections and derive the asymptotically stable equilibria of both complete and incomplete information games in a simple symmetric setting with two candidates. We also show how the model can be used to qualitatively explain several phenomena observed in reality: increases in costs of voting affect turnout adversely but there may be persistence of turnout levels between elections even though costs and other parameters change. Increase in uncertainty increases turnout while increases in the size of the electorate decrease it, in line with intuition.

JEL Classification: C72, D72

Keywords: Voter Participation, Voter Learning, Asymptotically Stable Equilibrium, Markov chain, long run equilibria

1 Introduction

The question of why people vote in large elections has occupied political scientists and economists since a long time (see e.g. Downs (1957)). The argument goes as follows: If people were rational and interested in voting in order to affect the outcome of the election (*instrumental voters*) then they would never vote. This is because voting is costly and the expected benefits small, since the probability that one vote counts is infinitesimal. However, observed turnout rates are quite high. This is the “paradox of voting” (Downs(1957)), Ferejohn and Fiorina (1974)).

The inability of existing models to explain high turnout has caused unhappiness with “rational choice” models of turnout which depend on the probability of being pivotal to provide incentives to vote (see for example Green and Shapiro (1996)). There is increasing emphasis on voting as a consumption decision, i.e. people vote because they like to – it gives them a “warm glow” (see Andreoni (2006) for a review of the literature on warm glow giving). These are the “expressive” theories of voting. Unfortunately, these theories at present are not very developed¹ and do not offer testable predictions. In particular, expressive theories cannot explain stylised facts such as increasing abstention rates over time as observed in many countries or the dependence of the turnout on the perceived distance between the electoral platforms of the candidates, or on the transactions costs of voting. These stylised facts seem to suggest that instrumental motivations do play a role, even if such motivations are not the only relevant ones in the voting decision. Indeed, empirical studies (Mueller, 1989, Ch.18) support the instrumental voter paradigm, although they often contradict each other in the details. On the other hand, instrumental voter theories may ask for too much rationality on the part of voters – being able to compute pivot probabilities does not seem a task we expect an ordinary voter to engage in before he makes his decision!

The main innovation of this paper is to adapt the standard instrumental voter model to more plausible boundedly rational voters who learn through a process of updating their behaviour in response to opinion polls. We build on the Palfrey and Rosenthal (1983,1985) (henceforth PR) model of turnout in a two candidate world and introduce learning in this framework. The learning process is based on voters reducing their turnout when the opinion polls give them information that implies that net benefits from voting are negative and vice versa.

The PR (1983) model has a plethora of equilibria², some with low turnout and some with high, and moreover some equilibria with counter-intuitive comparative statics. Imposing the learning process of

¹See however, Feddersen, Gailmard and Sandroni (2007) and Schuessler (2000) for theories that do provide some predictions

²We should remark here that the problem of multiplicity of Nash equilibria is not specific to the voting problem. The voting decision can be seen as a two sided version of the problem of private provision of a public good where the ‘public’ good is the fact of having one’s candidate winning and the contribution is the cost of voting. Indeed both cases present similar problems (free-riding, instability of the equilibrium in which everybody contributes as the number of players increases etc.). There are obvious co-ordination problems in both these settings (Bagnoli and Lipman (1989)). The methodology of this paper is thus equally applicable to the more general “team games” (PR, 1983).

our model, we are able to discard the unappealing equilibria of the PR model: indeed the very reason why these equilibria are at odds with intuition (the comparative statics) enables us to discard them using our learning process!³

We derive familiar and some new comparative statics results that yield testable predictions and are consistent with stylized facts about elections. The main results are: first we pinpoint the stable or stochastically stable equilibria in the basic model as well as the model with incomplete information. Second, we show the comparative statics with regard to these equilibria. When costs increase or the size of the electorate increases, abstention increases, in line with intuition; higher uncertainty about costs, in contrast causes turnout to increase; social norms and habits are important in voting decisions: if turnout levels were historically high, they tend to stay high; for example, voter turnout in Belgium, Australia, Italy is persistently high (close to 90% on average) compared to the United States (50% on average), there is persistence in the system. We then extend the model to repeated elections, which allows us to make an interesting prediction: for small electorates, sudden downward jumps in turnout may be observed; while for large electorates turnout changes smoothly with cost and benefit parameters. This corresponds to a phenomenon that is commonly observed in Western democracies: that of steadily falling turnout rates over time (<http://www.idea.int>). So why should an adaptive process be a good model of how voters behave? We think there are two types of reasons for believing this: methodological and empirical. We are studying an anonymous game with a large number of players who have little possibilities of communicating and who have little a priori information on what others will do. It is not reasonable to assume that they will guess the Nash equilibria or that they believe that everyone else will do so. This applies 'a fortiori' to all the selection criteria that are used to refine Nash equilibria (it is easy to see that these have no cutting edge in our problem anyway). It seems more sensible to assume that during the period that precedes the election players will collect the information that is available to them (e.g. statistics from previous elections, opinion polls etc.), and then find what is their optimal behaviour via some type of Gedanken Experiment.

It is interesting to see that the outcomes of such processes not only are Nash equilibria, but satisfy the strongest known rationality criteria (although the converse is not true – see Demichelis and Ritzberger (2003)). On the empirical side there is good evidence (Cooper et al (1990), van Huyck et al (1990)) that players do not apply a rationalistic selection criterion when they select strategies to play. Thus, for instance weakly dominated strategies play a role in their choice, and adaptive considerations seem to play a strong role in guiding their behaviour.

The framework of this paper is as follows: We discuss related literature in Section 2. We then introduce the model of voter learning in Section 3, then we consider the complete information game in Section 4 and the incomplete information game in Section 5. Section 6 has some discussion on how our model fares when faced with some stylised facts on elections and what we can say about expressive and

³We thank an anonymous referee for suggesting this nice justification.

instrumental voting. Finally Section 7 concludes.

2 Related Literature

There has been a lot of recent literature looking at the problem of how to explain empirically observed levels voter turnout (see Feddersen (2004) or Dhillon and Peralta (2002) for surveys). The first game theoretic investigations of instrumental voting were developed by Ledyard (1981) and Palfrey and Rosenthal (1983,1985).

PR (1983,1985) examined the issue of voter turnout in two different settings, one with complete information and one with incomplete information. They found that even in the simplest model with complete information (two candidates, two types of voters, symmetric equilibria) there was a problem of multiple equilibria. In particular, even with an equal number of the two types of voters in the population and restricting attention to symmetric mixed strategy equilibria⁴ they found two types of mixed strategy equilibria for plausible cost levels— one with low turnout and one with substantially high turnout. They say that this high turnout equilibrium “has the unappealing feature that there is another equilibrium with almost no one voting. Apparently the only reason the upper one (high turnout equilibrium) can be sustained is that the two electorates are of the same size so that for the probability of voting very close to 1, the probability of a tied election is very high. Again, the result rests on the fact that in equilibrium there is essentially no strategic uncertainty.” (PR, 1985). Of course, there is no reason why this argument should apply only to the high turnout equilibrium.⁵ They show that moving to the corresponding incomplete information game gets rid of this high turnout equilibrium. Indeed, in their model of the incomplete information game, they show that with some assumptions on the type of uncertainty allowed, the only symmetric mixed strategy equilibrium that survives is the low turnout one.

While we agree with PR (1985) that the low turnout equilibrium is intuitively more appealing, we argue that the reason they cite may not be the most compelling one. In fact, we show that removing some of the assumptions PR (1985) make on the type of voter uncertainty, the incomplete information model may still lead to multiple equilibria (at least for populations that are not “too” large). We show that considering the dynamics of reaching a Nash equilibrium by boundedly rational agents and incorporating in the model an important feature of real elections, the information aggregation devices like pre-election polls that exist, and that can co-ordinate people’s beliefs,⁶ would lead in a very simple and natural way to the low turnout equilibrium in their example.

⁴There may be objections to studying mixed equilibria, but we believe that they are quite justified in this setting – as psychological experiments have shown that choices between alternatives that are perceived as being similar tend to be relatively random (Fudenberg and Levine (1998) Chapter 4), besides the usual purification arguments.

⁵If we use the standard deviation as a measure of “strategic uncertainty”, one can see that it is the same for the two equilibria.

⁶See Fey (1997) for a similar idea but in the context of costless voting and multiple candidate elections.

Myerson (1998) adds another source of uncertainty in elections: while PR (1983, 1985) assume that the size of the population and the size of each group is known with certainty, Myerson (1998) considers games with population uncertainty. In a setting where the uncertainty is modeled by a Poisson process, it is shown that the low turnout equilibrium is the unique prediction of the modified PR model.

Bounded rationality models of voter turnout have been proposed by a number of authors. Among these, the closest to our paper is Diermeier and Van Mieghem (2009), who independently⁷ study the problem of turnout in large elections using a model of boundedly rational agents. Their model is quite different from ours: in each period a randomly selected voter receives information about turnout through polls and plays a best response based on Bayesian updating. With the “right” (not too much and not too little) amount of noise in polling data, they show the existence of a unique limiting distribution of the learning process. They show too that large turnout is possible even in large elections when the voting factions are of similar size.

There is also recent work in this category by Conley, Toossi and Wooders (2001), and Sieg and Schultz (1995). They are different from our model in that they are pure evolutionary models where some citizens are pre-programmed to turn out with a high probability, and they consider the conditions under which the strategy of high turnout survives. Benabou and Tirole (2006), on the other hand, study the social signalling incentives to take part in prosocial activities like voting. Funk (2005) found that the introduction of postal voting in Switzerland failed to increase turnout despite lower costs. This was attributed to the reduced benefits of showing that you did your civic duty (which signals cooperation and makes them more attractive as business partners).

We assume that information about candidates is complete and that there are conflicts between citizens on which candidate is best (private values). If we relax the assumption of complete information and private values, the turnout decision acquires another dimension – abstention may become strategic in order to let informed voters make the correct decision. In the context of costless voting this problem is studied by Feddersen and Pesendorfer (1996,1999).

An article on why turnout is related to the expected closeness of the election (and which does not rely on the probability of an individual voter being pivotal) is the model of Shachar and Nalebuff (1999). They model voters as passive followers who vote if they are contacted by a party. It is parties who make the strategic decision of how much effort to put into the campaign and who take account of how close the election is expected to be.

We do not study normative properties of turnout: in this paper we restrict attention to the PR (1983) model where the size of the two groups is the same. Thus the lower is turnout the more efficient it is, as the outcome is the same in all equilibria⁸. Borgers (2004) studies which voting rules are good when voting is costly.

⁷We thank an anonymous referee for pointing out this work to us.

⁸This feature of the model is discussed later in Section 7.

3 The Model

In this section we briefly describe and build on the model due to PR (1983, 1985).

There are two candidates (or two alternatives): 1 and 2. The voting rule is Simple Majority Rule: in case of a tie a coin toss takes place. There are N voters in the population. There are two groups of voters: T_1 (with N_1 voters who prefer 1) and T_2 (with N_2 voters who prefer 2) and all voters belong to one of these groups. Voting is costly and the cost of voting is the same for all voters. Voters have two pure strategies: vote (participate) or abstain – if they vote they always vote sincerely (i.e. for their best candidate). We ignore the strategy of voting for the candidate that is less preferred as this strategy is strictly dominated by abstention, and cannot be a part of any Nash equilibrium.

Let R represent the expected net benefit from voting, p the probability of being pivotal, C the cost of voting, B the difference between the benefits of i 's more preferred alternative winning as opposed to the less preferred one, and D a fixed benefit from the act of voting (civic duty). Let n_r^i , $r = 1, 2$ denote the number of voters excluding voter i who turn out to vote for candidate r . Let α_1 denote the probability that $n_1^i = n_2^i$, α_2 the probability that $n_1^i = n_2^i - 1$, α_3 the probability that voter i is not pivotal and candidate 1 leads, while α_4 denotes the probability that voter i is not pivotal and candidate 2 leads. Note that $\sum_k \alpha_k = 1$, $k = 1, 2, 3, 4$.

Thus the expected utility from voting for voter i is given by:

$$\alpha_1 B + \alpha_2 \frac{B}{2} + \alpha_3 B + \alpha_4 0 - C + D \quad (1)$$

The expected utility of voter i from abstaining is denoted by:

$$\alpha_1 \frac{B}{2} + \alpha_2 0 + \alpha_3 B + \alpha_4 0 \quad (2)$$

Therefore R is the difference between 1 and 2, and has the following expression:

$$R = (\alpha_1 + \alpha_2) \frac{B}{2} - C + D \quad (3)$$

Let $c = C - D$ represent the net cost of voting. We can normalise by dividing throughout by B so that only the ratio of costs to benefit matters. So, w.l.o.g set $B = 1$. Let $p = \alpha_1 + \alpha_2$. Thus a voter i will vote if and only if $p - 2c \geq 0$.

Let us first recall from PR (1983), the Nash equilibria of this game.

Let the probability for player i to choose to vote be q_i . Then (q_1, q_2, \dots, q_n) is a mixed strategy Nash equilibrium, if for all i voting and non-voting give the same expected payoff, given the mixed strategies of other players. As PR (1983) shows there are many Nash equilibria of this game, symmetric and asymmetric, pure strategy and mixed strategy.

We illustrate our model with the multiple equilibria that arise even when it is assumed that $N_1 = N_2$, and attention is restricted to symmetric equilibria. This case exhibits the same problems of multiplicity of equilibria as the general case, but allows us to make several technical simplifications in the solution that can be discussed at a very intuitive level.

The symmetric mixed strategy equilibrium in this game with $N_1 = N_2 = M$ is denoted q and it satisfies the following equation:

$$2c = \sum_{k=0, \dots, M-1} \frac{(M-1)!}{(M-1-k)!(k!)} \frac{(M)!}{(M-k)!(k!)} q^{2k} (1-q)^{2M-2k-1} \quad (4)$$

$$+ \sum_{k=0, \dots, M-1} \frac{(M-1)!}{(M-1-k)!(k!)} \frac{(M)!}{(M-1-k)!(k+1!)} q^{2k+1} (1-q)^{2M-2k-2}$$

If $0 < c < 1/2$ then this equation has either no solution, 1 solution or two. Figure 1 shows (as in PR, 1985) the equilibria. Note that to plot the graph and discuss the equilibria we use the conjecture made by PR (1983), reproduced below.

Conjecture 1 (PR, 1983) *For each integer $M > 1$, there exists a minimum voting cost denoted by c_{min}^M such that: (1) The unique symmetric equilibria at c_{min}^M is $q = 1/2$. (2) $c(q) = c(1-q) \geq c_{min}^M$ for $q \in (0, 1)$. (3) For every $c \in (0, 1/2)$ and every $\epsilon > 0$, $\exists M_\epsilon$ such that for all $M \geq M_\epsilon$ there exist exactly two equilibria, $q_H(c), q_L(c)$ such that $0 < q_L < \epsilon$ and $1 - \epsilon < q_H < 1$.*

In addition to these equilibria note that when $N_1 = N_2$ there always exists a pure strategy equilibrium, where all voters turn out to vote: this is because in this event the probability of being pivotal is 1, a tie is induced and the cost is less than $\frac{1}{2}$. The proof is by numerical calculation for small N . For large N , the conjecture is supported by a heuristic argument based on approximating the binomial with the normal distribution⁹. Let $g(q, M)$ denote the probability of being pivotal as a function of the equilibrium probability of voting of all voters. To see how the function $g(q, M)$ changes as M increases see Figure 4 in PR (1983).

⁹This is available upon request from the authors.




Figure 1

In Figure 1, we illustrate Conjecture (1). We have costs and the probability of being pivotal $p = g(q, M)$ on the horizontal axis. The vertical axis has q . Notice that when $q = 0$ or $q = 1$, then $p = 1$, and p is minimized when $q = \frac{1}{2}$. Hence the function $g(\cdot)$ is non-monotonic with respect to q : it is maximized at $q = 0, 1$ and minimized when uncertainty is the highest. Also, there is zero turnout if $c > \frac{1}{2}$, so in the figure we show costs between 0 and $\frac{1}{2}$. The points where $g(\cdot)$ intersects the vertical line showing the exogenously given cost c_0 , show the two mixed equilibria, q_H and q_L which are equidistant from the upper and lower horizontal lines, as suggested by part (2) of Conjecture (1). At the point where $q = 1$ there is always a pure strategy equilibrium, for any cost level, where everyone votes. This is denoted by FT (Full Turnout) in the figure. At the point where $q = \frac{1}{2}$, there is only one intersection (tangency) with the cost line, illustrating point (1) of the conjecture. This cost level is c_{min} and at this point the equilibrium turnout is $q = \frac{1}{2}$. Below this cost level there is only the pure strategy equilibrium FT . Figure 1 shows that as the cost increases beyond c_{min} there are two types of mixed strategy equilibria: the high turnout mixed equilibrium (denoted q_H) and one with almost no one voting (denoted q_L). In addition there is a pure strategy equilibrium, FT in which everybody votes. Let us analyze them separately: in the pure strategy equilibrium everybody goes to vote believing that everybody else will go. Such a situation may arise when the number of voters is small but seems

unlikely for electorates of large size. Note also that it relies heavily on complete information on the number of voters in each group¹⁰. The high turnout equilibrium q_H is counterintuitive, particularly if one considers comparative statics: it predicts an increase in turnout when the cost of voting increases. The low turnout equilibrium q_L , has good comparative statics properties. Note also that it Pareto dominates the other two. The outcome (50% win for each candidate) is the same in all three equilibria, even though turnout is different (in any symmetric *mixed strategy* equilibrium with identical costs of voters, the probability that a voter is pivotal is the same for all voters).

Finally as M increases to $M' > M$, the function $g(\cdot)$ shifts as shown to $g(q, M')$ so that c_{min} decreases: we can increase M so that q_L goes towards 0 and q_H goes to 1. This illustrates point (3) of the Conjecture (1).

PR (1985) state that the high turnout equilibrium is not robust in that they arise because of the fact that there is “essentially no strategic uncertainty” (see footnote 3) in the complete information model. We claim that this is not the main problem with the high turnout equilibrium. In fact the problem of multiple equilibria remains in the incomplete information model: it will be shown below (see Section 5) that at least in small populations, unless the amount of uncertainty is quite large (probably too large to be observed in concrete applications), not only does one still get three equilibria near the original ones but, in certain cases, many more may appear. Standard refinements like trembling hand perfection do not have any bite.

We claim, instead, that the main problem with the equilibria other than the low turnout one is that they are inconsistent with reasonable models of voter learning. Indeed, the missing piece of the model, opinion polls, and the trial and error process that arise quite naturally in such a setting is a good way to select the “robust” low turnout equilibrium, and this is what our model does. There is a very intuitive reason why the high turnout equilibrium is not a robust one, and this is true regardless of whether the game is of complete information or not. This is because of its stability properties with respect to learning dynamics. But first we will describe our model of voter learning.

4 Complete Information

In this section we will discuss learning in the PR model with complete information. Its primary purpose is to present the learning dynamics and show its use in selecting away “bad” equilibria. The artificial assumption of identical voting cost for everybody will be removed in section 5.

4.1 Dynamics: Single Elections

Learning can be modeled in several different ways (see Fudenberg and Levine (1998) for an overview of the literature). A good process of learning should satisfy the, possibly conflicting, requirements

¹⁰The assumption needed here is merely knowledge of the number of voters in each group, not equality: if $N_1 \leq N_2$ there is still a (mixed-pure) equilibrium in which everybody in group 1 votes based on similar beliefs.

of assuming as little rationality as possible from the players and of being efficient enough to give acceptable outcomes (we want the voters to learn something before election day comes...). Typical models are the best reply dynamics (as in the Cournot adjustment model), and its improvement (discrete) fictitious play (Brown (1951)). These assume rather myopic players but have the defect that sometimes they forecast players behaving in a rather unintuitive way, overshooting their beliefs and strategies so as to cycle for ever around an equilibrium without ever coordinating on it (Fudenberg and Kreps (1993)).

The problem can be circumvented, without having to assume a more sophisticated behaviour of the players, using a dynamic in the spirit of continuous fictitious play, a model in which they gradually adapt their strategy instead of jumping suddenly from one pure strategy to another (Demichelis and Germano (2002), Fudenberg and Levine (1998) Chapter 4). See also the justification for the case of repeated elections in Section (4.2).

As before, $g(q, M)$ denotes the probability of being pivotal as a function of the equilibrium probability of voting of all voters (computed above in the RHS of expression (4) with the coin toss tie breaking rule). Then, as before, the best reply is to vote if

$$g(q, M) - 2c > 0 \tag{5}$$

and abstain in the opposite case with equality corresponding to indifference between the two strategies.

We describe the model of learning now, assuming a fixed M (so we drop M as an argument in the function $g(\cdot)$ now). A voter starts by observing a q , e.g. the share of the population who voted on the last election or the result given by a poll (we assume M is not too small so that the empirical observation of this quantities is a reasonable estimator of q) and checks whether the inequality (5) tells her to go to vote (or to abstain). She realizes that everybody will do the same and so expects the correct q to be higher (lower). Once she has adjusted q she checks again what is her best reply and so on. We might think of this process as a repeated opinion poll. If this process converges to some point q^* , she will apply the corresponding mixed strategy.¹¹

The dynamics can be described by the differential equation¹² below, where the symbol M has been suppressed as the analysis is done for a fixed M . The state variable is q , the symmetric mixed strategy for all players. The state space is thus the space of mixed strategies, $X = [0, 1]$. The function $K(q)$ is the vector field.

$$\frac{dq}{dt} = K(q) \tag{6}$$

¹¹We may also think of voters using pure strategies and q describes the proportion of the population that votes – then our assumption is that population shares move continuously with changes in the expected benefit of voting. See Fudenberg and Levine (1998) Chapter 4 for more discussion on this point.

¹²If the adjustment steps are small enough, taking a discrete adjustment process would give essentially the same results

and sign $K(q) = \text{sign}(g(q) - 2c)$. The functional form of $K(q)$ depends on the particular speed of learning. Different individuals could have different K ; as long as the sign requirement on the possibly different $K_i(q)$ is respected it will not affect the results. So, to simplify notation, we assume that everyone adjusts in the same way. We also assume that the initial point is the same for all voters. Again, the hypothesis of a common initial q is not very restrictive—all of our results go through with small differences in the initial q . The assumption that voters are able to compute $g(q)$ exactly may seem at variance with the bounded rationality hypothesis, but for our purposes we could as well assume that $g(q)$ is some approximation of the true function, as long as it has the same shape.

We assume that the function $K(q)$ satisfies Lipschitz continuity and so there exists a unique solution for any initial q_0 , for all $t \in R$. The solution of this equation for any initial condition, $q(q_0, t)$ is continuous.¹³ Note that our requirements on the dynamics are very mild – any payoff consistent (and so any payoff monotonic or payoff positive (Weibull (1995), Demichelis and Ritzberger (2003))) dynamic satisfies it.

Now we examine the behaviour of the dynamics on the strategy space. The basic intuition is that the outcomes of the learning process that are likely to be seen in concrete cases are the limit points of $q(t)$ as t goes to infinity.

Before stating our results, we reproduce some definitions adapted from Weibull (1995): Let $q(q_0, t)$ be the solution of the differential equation (6) with initial condition $q(0) = q_0$. Let $q(q_0, t) \in X = [0, 1], \forall t \in R$.

Definition 2 *A state $q^* \in X$ is said to be Lyapunov stable if every neighborhood B of q^* contains a neighborhood B^0 of q such that $q(q_0, t) \in B$ for all $q_0 \in B^0 \cap X$ and $t \geq 0$.*

Intuitively a state is Lyapunov stable, or just stable, if no small perturbation away from it induces a movement away from it.

Definition 3 *A state q^* is asymptotically stable if it is Lyapunov stable and \exists a neighborhood B^* such that the following holds for all $q_0 \in B^* \cap X$:*

$$\text{Lim}_{t \rightarrow \infty} q(q_0, t) = q^* \tag{7}$$

A point that is not stable will be called unstable. While stability requires that there be no pull away from the state, asymptotic stability requires in addition that there be a local pull towards it as well.

Definition 4 *Basin of attraction of state q^* is the set of points $q_0 \in C : q(q_0, t)_{t \rightarrow \infty} \rightarrow q^*$.*

Intuitively, the basin of attraction of q^* is the set of initial conjectures $q_0 \in C$ that, with learning, will lead to q^* .

Now we can state Proposition (5):

¹³See Weibull (1995) Appendix.

Proposition 5 For any learning dynamics of type (6) the low turnout equilibrium q_L and the pure strategy full turnout equilibrium, FT , will be asymptotically stable, while the high turnout mixed equilibrium q_H will always be unstable.

We refer to the appendix for the (elementary) proof. Here is an informal discussion: If $c > \frac{1}{2}$ or $c < c_{min}$ any trajectory trivially converges to the unique equilibrium which is the zero turnout or the full turnout equilibria respectively. When $c_{min} < c < \frac{1}{2}$ the qualitative behaviour of the dynamics is as in Figure 2 below:

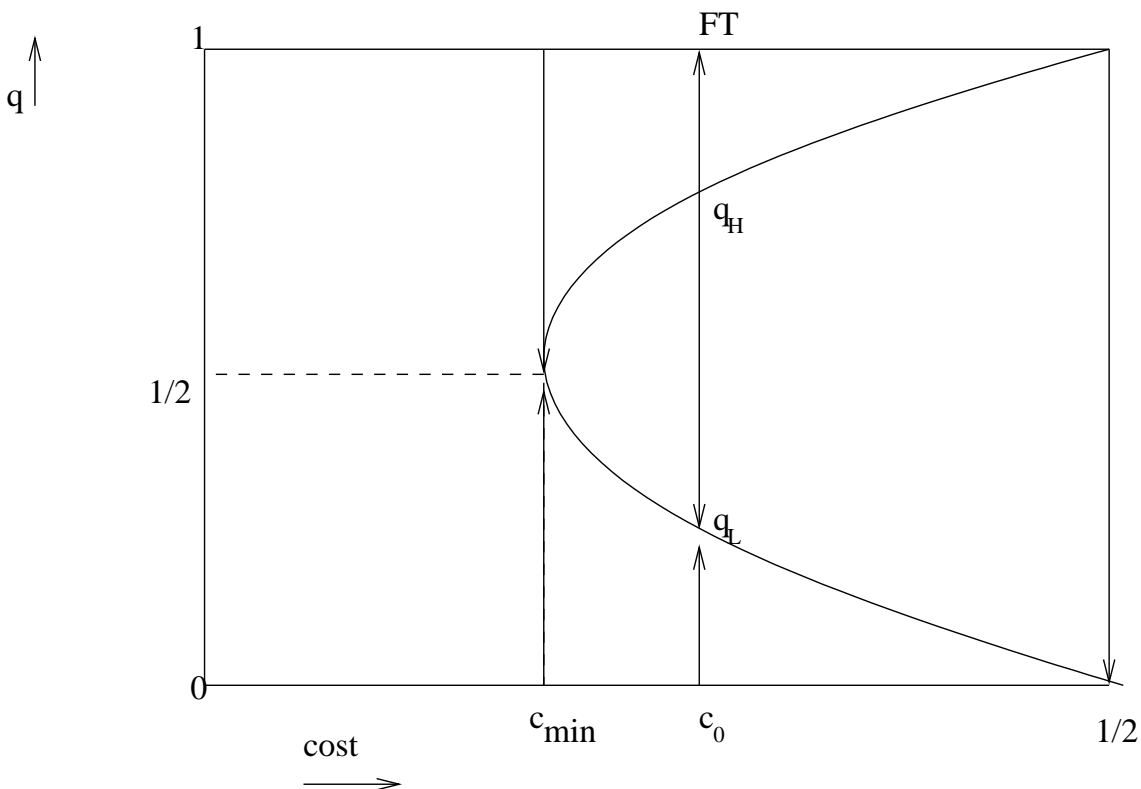


FIGURE 2

Any trajectory starting in the interval $[0, q_H)$ (its basin of attraction) converges to equilibrium q_L , therefore it is stable ; the high turnout mixed equilibrium q_H is unstable, no trajectory leads to it; equilibrium FT (the pure strategy equilibrium where everyone turns out) is stable with basin of attraction $(q_H, 1]$. Moreover, the basin of attraction of the low turnout mixed equilibrium q_L is larger than that of FT for any cost higher than c_{min} , and as c increases the basin of attraction for equilibrium FT shrinks.

Considerations about the size of the basin of attraction allow us to improve the predictions of the model: for a stable equilibrium to have a large basin of attraction means having many initial conditions leading to it and therefore means that it has a high probability of being observed. Note

also that we do not have to assume that players start the learning process at the same point, provided all the points are in the same basin of attraction. Assume that q_0 is distributed uniformly on $[0, 1]$ for any $c \in (0, \frac{1}{2})$. Then the probability that $q_0 \rightarrow q_L$ for any cost level $0 < c < \frac{1}{2}$ is given by $F(q_H) = q_H$ and the probability that $q_0 \rightarrow FT$ is given by $1 - F(q_H) = 1 - q_H$.

Then the prediction of the model in case of a single election is summarised in the following proposition.

Proposition 6 *For any learning dynamics of type (6), any integer M and any cost level c , such that $c_{min}^M < c < 1/2$, and q_0 uniformly distributed on $[0, 1]$, equilibrium q_H will never be observed, equilibrium q_L will be observed with probability q_H and equilibrium FT will be observed with probability $1 - q_H$. For every $c \in (0, 1/2)$ and every $\epsilon > 0$, $\exists M_\epsilon$ such that for all $M \geq M_\epsilon$ equilibrium q_L will be observed with probability $1 - \epsilon$ and equilibrium FT will be observed with probability ϵ . Moreover as $c \rightarrow \frac{1}{2}$, the probability of observing q_L goes to 1.*

The proof of this proposition is quite simple and is in the Appendix.

4.2 Repeated Elections

Suppose that elections are repeated regularly: we can index elections with $i = 0, 1, 2, \dots$. All other elements of the model remain the same as before. Now, at election i , voters begin the learning process at q_0^i . This depends on the turnout in the preceding election q_∞^{i-1} in the following (non-deterministic) way: q_0^i is a random variable uniformly distributed on the interval $[\underline{q}; \bar{q}]$ with $\underline{q} = \max\{0; q_\infty^{i-1} - \delta\}$ and $\bar{q} = \min\{1, q_\infty^{i-1} + \delta\}$; δ is a small positive number that gives a measure of the maximum possible mistake in ascertaining the turnout. Thus we get a random dynamical system, a Markov chain, since the system is time independent, whose states are the stable Nash equilibria, q_L and FT . For the following, we will denote the q_L equilibrium as 1 and the q_H equilibrium as 2, while the pure strategy equilibrium, FT is denoted by 3. The transition matrix for the Markov chain is denoted by P and is as follows:

$$\begin{bmatrix} p_{11} & p_{13} \\ p_{31} & p_{33} \end{bmatrix}$$

where p_{ij} = the probability that the state moves from state i to state j , with $i, j \in \{1, 3\}$. Thus the k step transition matrix is denoted P^k . The starting state probability vector $s(0)(M)$ (for a given size of the population M) is given by the size of the basins of attraction of the two stable equilibria. The behaviour is given by proposition (7):

Proposition 7 *For any $\delta > 0$, if the number of voters is larger than $N_0(\delta)$, the limit distribution of the outcomes q_∞^i is concentrated on equilibrium q_L .*

Proof in Appendix.

Note that the precise form of the probability distribution of the q_0^i is, to a large extent, irrelevant to the result, provided the distribution is concentrated around q_∞^{i-1} . Using a terminology borrowed from mechanics, equilibrium 3 is called "metastable". Intuitively the fact that the basin of attraction has positive measure but is very small means that if there are random disturbances, the equilibrium FT will be stable for a while but the probability of observing a jump to equilibrium q_L within time T goes to 1 when T goes to infinity. Equilibrium q_L is "stochastically stable" in the sense of Kandori, Mailath and Rob (1993). Proposition 7 can be generalised for other distributions. Suppose we have a distribution μ . Let α_{ij} denote the probability that q_0^i is in the basin of attraction of equilibrium j conditional on q_∞^{i-1} being in equilibrium i . Then the transition matrix $P(\mu, M)$ is given by:

$$\begin{bmatrix} \alpha_{11}(\mu, M) & \alpha_{13}(\mu, M) \\ \alpha_{31}(\mu, M) & \alpha_{33}(\mu, M) \end{bmatrix}$$

Let x_k, y_k denote the probability, after k iterations, of being in state 1 and 3 respectively. We have the following proposition:

Proposition 8 *For any initial probability distribution (x_0, y_0) , $(x_0, y_0)P_M^k$ converges to (x_M, y_M) independent of (x_0, y_0) , as $k \rightarrow \infty$. Moreover, as $M \rightarrow \infty$, we have $x_M \rightarrow 1$ and $y_M \rightarrow 0$.*

Proof in Appendix.

The intuition is that with a large number of voters and after a long enough period, the pattern most likely to be observed will be a long sequence of equilibrium q_L with transitions to equilibrium FT rare or completely absent.

This argument carries over to our discussion of the incomplete information case in Section 5. The only difference is that there may be several metastable equilibria instead of one, but the basin of attraction of equilibrium q_L is still the largest.

So, a consequence of metastability is that, if elections are repeated, equilibrium FT tends to jump to equilibrium q_L after a long sequence if the cost is below $\frac{1}{2}$ but higher than a certain value c_{min} . If the cost is not in this range there is only one equilibrium¹⁴.

4.3 Comparative Statics and Hysteresis

Note that, in all cases, equilibria with positive probability predict a nondecreasing q when c decreases, as intuition suggests. It is interesting to investigate further what happens when the cost, or the interest in the outcome, changes from one election to the other. In this case it is natural to ask that our version of the fictitious play dynamics starts at the percentage of voting of the last election. To

¹⁴Equilibrium FT tends eventually to jump to equilibrium q_L even if the number of voters M is small, the only change is in the average time it may take (e.g. if it is much more likely that people make small mistakes than larger ones, then for small M , if the initial state is FT , it remains FT , until a big enough mistake is made that puts voters in the basin of attraction of q_L); such a behaviour is typical of risk dominated equilibria in the sense of Harsanyi and Selten, 1988. See also discussion on public goods earlier.

simplify the discussion and isolate the different effects we will assume that there are no mistakes, i.e. that δ is zero. This allows us to see how q varies as a function of c . Suppose at some time we are given $c = c_0$ and the system is in equilibrium q_L . Then suppose that the introduction of electronic voting, for instance, causes c to decrease so that the predicted effect is a little increase of turnout (see Figure 2). But when c decreases below c_{min} , the mixed equilibria disappear and we suddenly fall in the basin of attraction of equilibrium FT , so q suddenly jumps up, i.e. at $c = c_{min}$ and below, equilibrium FT is the only stable equilibrium. The equilibrium will tend to persist for a while even if cost increases again until we reach the point where $c = 1/2$ and the only equilibrium possible is the pure strategy one, where nobody votes; alternatively even if c stays below $1/2$ but over c_{min} , in the "very long run" we will see q jumping down, because of the phenomenon of metastability described in the previous section. This is phenomenon is known as hysteresis or "memory" of the system, and explains why the same values of parameters can cause the emergence of different equilibria, depending on the initial state.¹⁵

Thus, we should observe phenomena of this type:

Corollary 9 *Suppose that $\delta = 0$ and $c < c_{min}$ then the only stable equilibrium is FT . The system persists at equilibrium FT even if cost increases as long as $c < \frac{1}{2}$. However, if $c > \frac{1}{2}$ then the equilibrium jumps to zero turnout.*

Moreover as M increases, for any fixed level of cost c , turnout decreases in the low turnout equilibrium q_L and remains the same in the equilibrium FT .

5 Incomplete Information

We capture uncertainty by a model of incomplete information about costs, exactly as in the PR (1985) model¹⁶. Each voter i has a cost of voting c_i which is private information to him. Let the cumulative distribution of costs be denoted $F(c)$ and for simplicity we assume the distribution to be the same between the two groups. We look for the Bayesian Nash equilibria of this game (as in PR (1985)). Each voter then has a decision rule that specifies whether to vote or not as a function of his own cost c_i . It is easy to see that in any symmetric Bayesian equilibrium a voter votes if his cost is below a certain threshold level, c^* , so in this case a learning process will be a dynamic on the c^* that is chosen. Thus a (symmetric) Bayesian equilibrium is a cost level c^* such that $2c^* = g(q(c^*))$, with the corresponding $q^* = F(c^*)$. This corresponds to the equilibrium outcome in the game of complete information where all voters have cost c^* and vote with probability $q^* = F(c^*)$. Note that, although it is natural to assume that players choose the cost level c^* , this is equivalent, for symmetric equilibria, to choosing q^* . All dynamics in terms of one variable can be easily translated in dynamics in terms of the other. So let $C(q)$ represent the inverse of $q(c) = F(c)$. Then we need $2C(q^*) = g(q^*)$. In

¹⁵We thank Jonathan Cave for pointing out this interesting feature of the model.

¹⁶PR (1985) also have a section on incomplete information about preferences but this has similar results.

the graph below (Figure 3), this equilibrium is given by the point where the distribution function intersects the curve $g(q)$, which shows the probability of being pivotal (as in Figures 1 and 2 above).




FIGURE 3

PR (1985) use assumptions under which there is only one intersection and the intersection converges to the point $q = 0$ as $N = 2M$ becomes large. As is evident from the graph, however, everything depends on the shape of the distribution function $F(c)$ (or its inverse $C(q)$).

We argue first that incomplete information does not by itself resolve the problem of multiple equilibria. The assumptions used by PR(1985) are: (1) $F(c)$ is continuous on $(-\infty, \infty)$, (2) $F(0) > 0$ and (3) $F(1) < 1$.

The first assumption is rather natural and corresponds to assuming that the probability distribution of c has no atoms. Assumption 2 is quite realistic also, i.e. that there is a positive probability that cost will be negative (civic sense will prompt some people to vote regardless of their assessment about being pivotal.) However, the last assumption is stronger: it implies that there is a positive probability that a voter would not show up even if he were sure to be pivotal. It is also not innocuous and PR (1985) have an example where relaxing this assumption takes us back to the problem of multiple equilibria in the complete information case (see Figure 3). Moreover, there seems to be an implicit assumption that $F(c)$ is not too wiggly: Figure 4 shows a case where the curvature of $F(c)$ can change quite fast so that many additional equilibria are introduced. Note that such a multimodal probability distribution is not so pathological: it could model a population made of different groups each with different costs and with small variance within a group. Thus, incomplete information does not solve the problem of multiple equilibria, sometimes it even introduces more of them, some of which have high turnout and our intuition suggests that they should be “non-robust” in some sense. However we

should mention that for it is *possible*, that sufficiently large M (or N), the PR (1985) conclusion is still valid. This however depends on the shape of the distribution function. To address this we now consider learning as before: $\frac{dq}{dt} = K(g(q) - 2C(q))$. We can isolate the stable equilibria as being the ones where $C(q)$ intersects $g(q)$ from below. These are the equilibria 1 (q_L) and 3 in Figure 4.




FIGURE 4

More formally we state this in the next Proposition: Let the graphs of $C(q)$ and $g(q)$ be in a generic position: this means that they intersect transversely i.e. $K'(0) \neq 0$.

Proposition 10 : *Assume $K'(0) \neq 0$. Then the asymptotically stable points are those such that $\frac{d(2C(q)-g(q))}{dq} \geq 0$.*

The proof is the same as for Proposition (5).

In the same way as in proposition (6), it is not hard to see that if M is large the only equilibrium with a large basin of attraction, containing at least the interval $[0, 1/2]$, is q_L – the low turnout equilibrium. All the others are either unstable, i.e. with zero probability, or metastable, i.e. with a small basin of attraction, whose size goes to zero when N goes to infinity.

Finally we remark that the analysis of repeated elections is the same as in Section 4, all results carry over. Thus, equilibrium q_L has the largest basin of attraction and will be the stochastically stable equilibrium whenever there is more than one equilibrium.

In the next section we do a similar analysis of the behaviour of equilibria when the parameters of the model are allowed to vary. We also illustrate with an example that the conditions under which

there is a unique equilibrium with incomplete information may be very restrictive.

5.1 An Example

We now investigate how equilibria change when the distribution of c varies in a simple example. We shall assume that the c are uniformly distributed on the interval $[\bar{c} - s; \bar{c} + s]$, so that we have two parameters the average \bar{c} and a measure of the dispersion s . Hence $F(c) = C(q)$. Other distributions, such as the Gaussian, can be discussed in the same way and give qualitatively similar results (see the discussion at the end of the section). A uniform distribution corresponds to an F whose graph is shaped as in Figure 5 below:




Figure 5

The slope is given by $\frac{1}{2s}$ and the intercepts are given by the point $\bar{c} - s$ and $\bar{c} + s$. The interpretation is that there are no atoms at any cost below $\bar{c} - s$. If, on the other hand there are people in the population with negative costs then the $F(c)$ line will intersect the vertical axis. So as \bar{c} increases, holding s constant the $F(c)$ line shifts to the right, parallel to itself. As s increases, on the other hand the $F(c)$ line shifts to the left and becomes flatter.

Now using the equilibrium result that $q^* = F(c^*)$, we have that $C(q^*) = F^{-1}(c^*)$. Now we can plot the inverse function against q and superimpose the function $g(q, M)$. In equilibrium, $g(q, M) = 2C(q)$, so the intersections between $2C(q)$ and $g(q, M)$ give the equilibria. This is shown in Figure 6.




Figure 6

Notice that the slope of the $g(q, M)$ line evaluated at $q = 1$ is $M - 1$ using the binomial formula from equation (5).¹⁷

Consider Figure 7. It should be geometrically intuitive, and it is easily proved, that if s is large enough so that the slope of the line GH ($\frac{1}{2s}$) $\frac{2}{M-1}$ is smaller than the slope of the arc BD (i.e. $\frac{2}{M-1}$) there is only one equilibrium. Since $\frac{2}{M-1} \geq \frac{1}{2s}$ iff $s \geq \frac{M-1}{4}$, it follows that:

¹⁷Recall that $N = 2M$.




Figure 7

Proposition 11 *If $s \geq \frac{M-1}{4}$, then there exists a unique equilibrium for any value of the average cost \bar{c} . In this range of s the unique equilibrium corresponds to nobody voting if $\bar{c} - s > 1/2$, everybody voting if $\bar{c} + s < 1/2$ and a percentage of voters that is a smoothly decreasing function of \bar{c} in the cases in between.*

The proof is obvious.

This corresponds to the case studied by PR (1985). Note that this is a very large value for s , for plausible values of M . The more realistic case of small (or rather not enormous), s , i.e. $s < 4(M-1)$, is more interesting and presents several analogies with the case of complete information, that corresponds to $s = 0$.

For any s , let $c(s)$ be the value of \bar{c} such that the line GH is tangent to the curve BD . Note that $c(s)$ is equal to c_{min} for $s = 0$ and increases monotonically to $1/2$ when $s = (M-1)/4$. It is also clear from the picture that $\bar{c} \leq 1/2 - s$. We now have, as in the complete information case, three cases:

Proposition 12 *Suppose that $s < (M-1)/4$. If $\bar{c} < c(s)$ there is one equilibrium with $q = 1$ (everybody votes) (shown by area A in Figure 8). If $c(s) < \bar{c} < 1/2 - s$ there are three equilibria: $q = 1$ (FT) which is stable with a small basin of attraction (which decreases with $c(s)$), an unstable one with large q (q_H) and a stable one with a low q (q_L) and a large basin of attraction. This corresponds to area B in Figure 8. If $\bar{c} + s > 1/2$, there is only one equilibrium with $q_L < 1$, and q smoothly decreasing with \bar{c} (shown by area C in figure 8).*

Corollary 13 *If the average cost, \bar{c} is fixed, and s , the measure of uncertainty on cost, increases, the stable equilibrium q_L has increasing turnout.*

This gives a very interesting insight into the relation between uncertainty and turnout, which has not been investigated in previous literature. The domain in the \bar{c}, s plane corresponding to the three cases is shown in Figure 6.




Figure 8

Most of our discussion applies to other distributions as well provided they are regular enough. Note however that the tangency between the F curve and the q curve may not be unique in particular cases, such as costs concentrated around a finite number of values. This would make the jump in the hysteresis cycle split into the composition of several smaller ones.

6 What can be explained?

The honest answer to this question would be: not everything. However, the qualitative predictions of our model accord well with the facts and provide reasonable interpretations of them. This is rather surprising given that the PR model is an oversimplified representation of real elections.

As mentioned before, one of our first comparative static results is that turnout is decreasing with costs and with the size of the population. Moreover, as we discussed in Section 5 we are able to predict a positive correlation between turnout and uncertainty, and this effect can be tested empirically.

Our model, and in particular hysteresis, can be used to see what happens when the cost c is changed by the introduction or removal of voting laws. For example: abstention is high in the U.S.A. and has been significantly lower in countries such as Belgium and Italy, even though there is no reason to expect big differences in cost of voting or interest in the elections. An explanation of this fact might be as follows: Belgium and Australia e.g. and in the past Italy had laws against abstention that made c quite low and possibly negative, so equilibria (are) were high turnout equilibria. The abolition, or lack of enforcement, of such laws has moved the state to the segment of higher cost but with persistence of large q ; in U.S.A. where there have never been such laws the more stable low turnout equilibria are observed.

A stylised fact about elections in Western democracies is that turnout rates have been falling steadily over time (<http://www.idea.int>). There are two possible explanations: one is the convergence of party platforms that leads to increasing apathy on the part of voters. An alternative explanation is that voters learn that there is another (low turnout) more efficient equilibrium and move towards that—see Section (3) in which we mention that the low turnout equilibrium Pareto dominates the other two, because the outcome is the same (tie between the candidates), but the smallest number of people incur the cost in the low turnout equilibrium. In our model the first explanation would correspond to an increase in the cost of voting (recall that this is a ratio of actual costs to benefits), while the second would correspond to a jump from equilibrium FT to equilibrium q_L described in our section on repeated elections. It is possible to distinguish between the two explanations, as our comparative statics predict a gradual decline in the first case but a sudden jump in the second case. Since the decline is indeed gradual with small oscillations, we may deduce that the first explanation is the correct one. The small oscillations would correspond to small changes in cost of voting and its distribution.

7 Conclusions

In this paper we showed how considerations based on learning dynamics and stability can be used to provide a realistic and plausible model of voter learning when faced with the turnout decision. The stable equilibria have good comparative static properties and give the model some predictive value. Incidentally, the model shows how to select the intuitive equilibria in the PR (1985) model without having to resort to ad hoc arguments.

Clearly this stylised model misses many features of real elections: usually there are more parties, voters are not homogenous – they have different costs and benefits, and there exist many other ways to influence policy than voting, such as lobby groups and other forms of direct action. Moreover

parameters such as the cost of voting or the measure of the interest in a candidate are not directly measurable with reasonable confidence, so that it seems hard to go beyond the rough estimates we made in the last section.

It should be obvious that in such cases asking for precise quantitative predictions is difficult – there are too many factors that can affect the predicted turnout that are hard to quantify exactly. So what can we hope to achieve from such simple models? The essence of the instrumental voter models is that voters are strategic – the probability of being pivotal matters to them, as do costs and benefits of voting. Thus the power of the model comes from predictions about comparative statics and other qualitative predictions. Since our model gives some sharp qualitative predictions (jumps, hysteresis, long time drifting away from metastable equilibria) that are very robust with respect to the parameters involved, it makes the model more apt to be tested in this way.

One objection to this approach that we anticipate is that the model is not really suited to large populations because for large populations the probability of being pivotal is very small, and hence the game theoretic model is more suited to smaller groups e.g. committees. A partial answer has been provided above, anyway our model is quite applicable to voting in committees as well, except for the assumption that voters are myopic in their learning behaviour. We plan to extend this model allowing for more sophisticated learning on the part of voters.

References

1. Andreoni, J. (2006), "Philanthropy", in S.C Kolm and J. Mercier Ythier, eds. *Handbook of Giving, Reciprocity and Altruism*. Amsterdam: North Holland.
2. Bagnoli and Lipman, 1989, "Provision of Public Goods : Fully implementing the Core through Private Contributions", *Review of Economic Studies*, Vol 56, pp. 583-601.
3. Beck, N., 1975, "A Note on the Probability of a Tied Election", *Public Choice* 23, Fall, pp. 75-80.
4. Benabou, R. and J. Tirole, 2006, "Incentives and Prosocial Behaviour", *American Economic Review*, Vol. 96, No.5, pp. 1652-1678.
5. Borgers, T., 2004, "Costly Voting", *American Economic Review* 94, 57 - 66.
6. Brown, G.W., 1951, "Iterative solutions of games by fictitious play", in "Activity analysis of production and allocation", ed. T.C. Koopmans, New York: Wiley.
7. Conley, J., Toossi, A. and M.Wooders, 2006, "Memetics and Voting: how nature makes us public spirited", *International Journal of Game Theory*, Vol.35 (1), December, pp. 71-90.
8. Cooper, R.W., D. DeJong, R. Forsythe, T.W.Ross, 1990, "Selection criteria in coordination games: some experimental results", *American Economic Review*, March 1990, 80, pp.218-233.
9. Demichelis, S. and F.Germano, 2002, "On knots and dynamics in games", *Games and Economic Behaviour*, Vol.41 (1), pp. 46-60.
10. Demichelis, S. and Klaus Ritzberger, 2003, "From evolutionary to Strategic Stability", *Journal of Economic Theory*, Vol. 113(1), pp. 51-75.
11. Dhillon, A. and S.Peralta, 2002 "Economic Theories of Voter Turnout", *The Economic Journal*, Vol.112 No.480, pp.F332-F352.
12. Diermeier, D. and Jan A.Van Mieghem, 2009 "Coordination and Turnout in Large Elections", *Mathematical and Computer Modelling* (forthcoming).
13. Downs, A., 1957, "An Economic Theory of Democracy", Harper, New York.
14. Feddersen, T. J. and W. Pesendorfer, 1996, "The swing voter's curse", *The American Economic Review* 86, pp. 408-424.
15. Feddersen, T. J. and W. Pesendorfer, 1999, "Abstention in Elections with asymmetric information and diverse preferences", *American Political Science Review*, vol.93(2), June, pp.381-98.

16. Feddersen, T., S. Gailmard and A.Sandroni, 2009, "Moral Bias in Large Elections: Theory and Experimental Evidence", *American Political Science Review*, Vol. 132(2), pp. 175–192.
17. Feddersen, T., 2004, "Rational Choice Theory and the Paradox of not voting", *Journal of Economic Perspectives* 18(1), pp. 99-112.
18. Ferejohn, A. F. and M. P. Fiorina, 1974, "The paradox of not voting: a decision theoretic analysis", *American Political Science Review* 68, pp 525–536.
19. Fey, M. 1997, "Stability and Coordination in Duverger's Law: A Formal Model of Preelection Polls and Strategic Voting", *American Political Science Review*.
20. Fischer, A.J., 1999, "The Probability of being Decisive", *Public Choice* 101, pp. 267-283.
21. Fudenberg, D. and D. Kreps, 1993, "Learning Mixed Equilibria", *Games and Economic Behaviour*, 5, pp.320-367 .
22. Fudenberg, D and David K.Levine, 1998 "The theory of learning in Games", MIT Press.
23. Funk, P., 2008, "Social Incentives and Voter Turnout: Evidence from the Swiss Mail Ballot System", *Journal of the European Economic association* (forthcoming).
24. Good, I.J. and Lawrence S.Mayer, 1975, "Estimating the Efficacy of a Vote", *Behavioural Science*, 20(1), January, pp.25-33.
25. Green, D. and I.Shapiro, 1996, "Pathologies of Rational Choice Theory: A critique of applications in political science". new Haven: Yale University Press.
26. Harsanyi, J. and R. Selten, 1988, "A General Theory of Equilibrium Selection in Games", MIT Press.
27. Kandori, M., G.J.Mailath, and R.Rob 1993, Learning, Mutations and Long Run Equilibria in Games, *Econometrica*, Vol.61, No.1, pp.29–56.
28. Kao, Edward P.C., 1996, *An introduction to Stochastic Processes*, Duxbury Press.
29. Ledyard, J. O., 1981, "The paradox of not voting and candidate competition: a general equilibrium analysis", in G. Horwich and J. Quirck (eds.), *Essays in contemporary fields of economics*, Purdue.
30. Mueller, Dennis C., 1989, "Public Choice II", CUP.
31. Mulligan, Casey B. and C.Hunter, 2003, "The Empirical Frequency of a Pivotal Vote", *Public Choice*, Vol. 116(1-2), pp. 31-54.

32. Myerson, R., 1998, "Population uncertainty and Poisson games", *International Journal of Game Theory* 27, 375–392.
33. Palfrey, T.R. and H. Rosenthal, 1983, A strategic calculus of voting, *Public Choice*, 41, pp.7–53.
34. Palfrey, T.R. and H. Rosenthal, 1985, Voter Participation and Strategic Uncertainty, *The American Political Science Review*, Vol.79, pp.62–78.
35. Shachar, R. and B. Nalebuff 1999, "Follow the Leader: Theory and Evidence on Political Participation", *American Economic Review*, Vol.89, No.3, pp.525–549.
36. Schuessler, A.A., 2000, *A Logic of Expressive Choice*, Princeton, Princeton University Press.
37. Sieg, G. and C.Schulz, 1995, "Evolutionary Dynamics in the voting game", *Public Choice*, vol.85, pp. 157–72.
38. van Huyck, R.B, R.C. Battalio, R.O.Beil, 1990, " Tacit Coordination games, Strategic uncertainty and Coordination failure", *American Economic Review*, March 1990, 80, pp. 234-248
39. Weibull, J., 1995, "Evolutionary Game Theory", MIT Press.

Appendix

Proof of Proposition 5:

Proof. Let q_L and q_H represent the low and high turnout mixed equilibria respectively. We show that any trajectory beginning in the interval $[0, q_H]$ converges to q_L , thus $[0, q_H]$ is the required neighbourhood B^* for equilibrium q_L , and any trajectory beginning in $(q_H, 1]$ converges to equilibrium FT , i.e. to $q = 1$, hence the required neighbourhood for equilibrium FT is $(q_H, 1]$.

Consider first a path starting in the interval $[0, q_H)$, i.e. $q(q_0, t) \in [0, q_H)$. By equation (6), $q(\cdot, t)$ is an increasing continuous function of t in this part of the domain, and remains so upto q_L . Hence $q(\cdot)$ must converge to q_L . The other direction is the same. ■

Proof of Proposition 6:

Proof. The basin of attraction for equilibrium $q_L(c)$ (or equilibrium 1) is given by $[0, q_H(c))$ and the basin of attraction for the equilibrium FT for fixed c is given by $(q_H(c), 1]$. Thus if $c < 1/2$, $q_H(c) > 1/2$, hence the basin of attraction of equilibrium of $q_L(c)$, i.e. q_H is bigger than $1/2$, while that of equilibrium FT $1 - q_H$ is smaller than $1/2$. As M increases moreover, by Conjecture 1 above, we have q_H tending to 1 for any fixed cost level c . Similarly as $c \rightarrow \frac{1}{2}$, the basin of attraction of equilibrium FT goes to zero and the basin of attraction for the low turnout equilibrium, q_L , goes to 1. ■

Proof of Proposition 7:

Proof. Let p_{ij} denote the probability of starting in state i and reaching state j . The transition matrix of the Markov Chain is given by:

$$\begin{bmatrix} p_{11} & p_{13} \\ p_{31} & p_{33} \end{bmatrix}$$

Since any point in $[0, q_H)$ converges to equilibrium 1 and any point in $(q_H, 1]$ converges to equilibrium 3, p_{11} is the probability that $q_0^i \in [0, q_H)$, the basin of attraction for equilibrium 1, if $q_\infty^{i-1} = q_1$. p_{13} is the probability that $q_0^i \in (q_H, 1]$, the basin of attraction of equilibrium 3 if $q_\infty^{i-1} = q_1$. p_{31} is the probability that $q_0^i \in [0, q_H)$, the basin of attraction for equilibrium 1, when $q_\infty^{i-1} = 3$ and p_{33} is the probability that $q_0^i \in (q_H, 1]$, the basin of attraction for equilibrium 3, when $q_\infty^{i-1} = 3$.

Note that when $N \rightarrow \infty$, $q_H \rightarrow 1$. Suppose that δ is small enough that $q_L + \delta \in [0, q_H)$, if $q_\infty^{i-1} = q_1$, the probability that it stays in state 1 is 1 (recall q_0^i is uniformly distributed in the interval $[q, q_1 + \delta]$ where $q = \min(0, q_1 - \delta)$); thus $p_{11} = 1$, and $p_{13} = 0$. However if the system starts from state 3, it may stay in state 3, if $1 - \delta < q_H$. However, by Conjecture 1 (PR, 1985), there is an $N_0(\delta)$ large enough (hence $q_H = q_0(\delta)$ large enough) that $1 - \delta > q_0(\delta)$. Thus the probability $p_{31} = F(q_0(\delta)) = \frac{q_0(\delta) - (1 - \delta)}{\delta}$, which is strictly positive. Of course $p_{33} > 0$.

Suppose now that δ is large so that $p_{11} < 1$. Then, by the above reasoning $p_{31} > 0, p_{33} > 0$, and there is an $N_0(\delta)$ large enough that $p_{11} = 1, p_{13} = 0$. Let $p_{33} = \beta > 0$. Then P^k has the following form:

$$\begin{bmatrix} 1 & 0 \\ 1 - \beta^k & \beta^k \end{bmatrix}$$

This clearly converges to

$$\begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}$$

Hence the limit distribution of the outcomes $q_\infty^i = (1, 0)$.

■

Proof of Proposition 8:

Proof. The first part of the proposition is an elementary fact of the theory of Markov Chains (e.g. Kao, 1996, Chapter 4). For the second part, note that (x_N, y_N) is the unique solution to the equations:

$$x_N + y_N = 1 \tag{8}$$

$$(1 - \alpha_{11}(\mu, N))x_N = (1 - \alpha_{33}(\mu, N))y_N \tag{9}$$

When $N \rightarrow \infty$ the Lebesgue measure of the basin of attraction of equilibrium 3 goes to zero, so by the assumption of absolute continuity, $\alpha_{33}(\mu, N) \rightarrow 0$, and $\alpha_{11}(\mu, N) \rightarrow 1$, so the claim follows.

■