Epistemic Solidarity as a Political Strategy*

ROBERT E. GOODIN

School of Philosophy
Australian National University
Canberra ACT 0200, Australia
<Bob.Goodinb@anu.edu.au>

Department of Government
University of Essex
Colchester CO4 3SQ, UK
<regoodin@essex.ac.uk>

KAI SPIEKERMANN

Corresponding Author

Department of Government
London School of Economics & Political Science
Houghton Street
London WC2A 2AE, UK
<k.spiekermann@lse.ac.uk>

Abstract

Solidarity is supposed to facilitate collective action. We argue that it can also help overcome false consciousness. Groups practice "epistemic solidarity" if they pool information about what is in their true interest and how to vote accordingly. The more numerous "Masses" can in this way overcome the "Elites," but only if they are minimally confident with whom they share the same interests and only if they are (perhaps only just) better-than-random in voting for the alternative that promotes their interests. Being more cohesive and more competent than the Masses, the "Elites" can employ the same strategy perhaps all the more effectively. But so long as the "Masses" practice "epistemic solidarity" they will almost always win, whether or not the "Elites" do. By enriching the traditional framework of the Condorcet Jury Theorem with group-specific standards of correctness, we investigate how groups can organize to support the alternatives truly in their interests.

Keywords: Condorcet jury theorem, information pooling, epistemic democracy, solidarity, false consciousness

* An earlier version of this paper was presented to a workshop at the Public Choice Research Centre, University of Turku and at a workshop on legitimacy and factual disagreement at the University of Copenhagen. We are grateful for comments, then and later, from Christian List, Dennis Mueller, Don Saari, Katri Seiberg, Martin Marchman Andersen, Klemens Kappel, Jurgen de Wispelaere, Theresa Scavenius and others.

Epistemic Solidarity as a Political Strategy

How can the Masses overcome the power of Elites? By organizing, of course. We have known for ages, and it is true in ever so many ways. The purpose of this paper is to draw attention to yet another, perhaps more surprising respect in which that is also true: organizing can be a way of overcoming a certain sort of false consciousness itself.

Traditional organizing aims at producing concerted action. In the present application, the aim of organizing is to produce correct beliefs – specifically, correct beliefs about our true interests. Traditionally, false perceptions of our interests are seen as an impediment to collective action. In the present context, awareness that our perceptions of our interests may be false serves as an impetus for us to collectivize – specifically, to pool our information.

Elites are advantaged in ever so many ways. In addition to having more power and wealth, they also have more information – most especially, information about what is in their interests and how to promote them. By organizing, the Masses can overcome those advantages. They do so by pooling, not only their power and wealth (in all the traditional ways), but also (the novel suggestion of this paper) their information about what is in their interests and how to promote them.

We dub this strategy of pooling information with selected others 'epistemic solidarity'. The strategy works only within limits. First, people have to be relatively confident *with whom* they share the same interests, even if they are unsure exactly *what* those interests are. Second, the people in the group thus identified have to be more likely to be *right* than random regarding the content of those interests.[1] Third, the less competent Masses must be more *numerous* than the more competent Elites. How much is required in each dimension is a function of how much is present in both of the other dimensions.

What makes the trick of epistemic solidarity work is a phenomenon familiar from discussions of the Condorcet Jury Theorem (CJT). That theorem says, roughly, that a majority among a group of voters, each of whom is more likely to be right than random, is more likely to be correct than is the individual voter; and the larger the number of voters, the more likely is a correct majority vote (that probability approaching certainty as the number of voters approaches infinity).

Here is one way of stating the theorem more precisely.[2] Assume a decision between two alternatives and a majority decision (without abstentions) in a population of odd size *N*. Let the state (of the world) be the fact which of the two alternatives is objectively correct. Two core assumptions are necessary for Condorcet's jury theorem to hold:

[1] Both of which are to say: false consciousness must not run too deep.
[2] Cf. Grofman, Owen and Feld 1983; List and Goodin 2001.

Competence. All voters have the same probability $1>p>\frac{1}{2}$ to vote for the correct alternative (and this is true for both states).

Independence. The votes are independent, conditional on the state.

The theorem can then be stated as follows:

Condorcet Jury Theorem. Given Competence and Independence, the probability of a correct majority decision increases in (odd) group size and approaches 1 as N goes to infinity.

The present application alters that traditional framework by respecifying what voters are right *about*, that is, the state. In the classic framework, the state is taken to be some truth about the world that is the same for everyone (how many jellybeans there are in the jar, or whether kissing transmits HIV, or what is in 'the common good' for us all). In the current application, we abjure notions of 'the common good' and focus

instead upon *group-specific* criteria of 'what is good for us', which differs from one group to another. This means that the state is now group-specific.[3]

Still, the same CJT logic applies – only for each group separately. Provided that the standard CJT conditions hold, a majority vote among members of each group is more likely to be correct than is an individual member about what is in her and her group's interest. Furthermore, that effect will be stronger the larger the group – which is of course precisely the advantage that the Masses enjoy over the Elites in availing themselves of this strategy.

The upshot of this paper will be that the Masses can pretty well count on winning, just so long as they practice epistemic solidarity and they have sufficiently independent and competent opinions to pool. There are some settings in which that will not be true, despite independence and competence.[4] But these settings are sufficiently extreme to be of little practical consequence.

Epistemic solidarity is a game that two can play, however. Furthermore, the Elites might well be better at playing it than the Masses. If the Elites succeed in practicing epistemic solidarity and the Masses do not[5] then smaller and individually more competent Elites can sometimes prevail

[3] The idea of a group-specific truth goes back to Alvin Goldman (1999, ch. 10). For a theorem in that regard, see ***.

[4] Where Elites are almost as big as Masses, for example, and/or are vastly more competent than them (while Mass competence is just over random).

[5] Or do so only very badly. For how bad the Masses have to be at pooling, see Sec. IV.A below.

over Masses who are more numerous but individually less competent. While that outcome will not always occur, it will in some scenarios that are sufficiently credible to be a real cause for concern.

I. The Effects of Sheer Numbers Alone

To some extent, the Masses can win through sheer force of numbers alone, even without practicing epistemic solidarity. They can afford more of their own to vote incorrectly, precisely because they have numbers to spare. Let us start by investigating the chances of the Masses winning in that baseline case, without any epistemic solidarity.

Suppose there are two groups in society, the Elites and the Masses. Suppose that the policy that is truly in the interests of each member of the Elites is $\mathscr{E}$ and of each member of the Masses is $\mathscr{M}$, and those are the only two options. Suppose that there are E Elite voters, each of whom is p_e likely to vote correctly from his point of view (i.e., for $\mathscr{E}$); and suppose that there are M Mass voters, each of whom is p_m likely to vote correctly from his point of view (i.e., for $\mathscr{M}$). The total size of the population is $N = E + M$.

Imagine now a direct referendum in which each voter votes sincerely and independently of one another (conditional on the correct answer for their group). And suppose that the electorate, as well as both subgroups E and M,

are large, so that the 'law of large numbers' applies. Then, as the population size goes to infinity while keeping the ratio E:M fixed, the proportion of votes for $\mathcal{E}$ in the total population would approach the population proportion of Elite voters who vote correctly from their point of view (which is approximately p_eE/N) plus the proportion of Mass voters who vote incorrectly from their point of view (which is approximately $(1-p_m)M/N$). The proportion of votes for $\mathcal{M}$ in the population would be the population proportion of Mass voters who vote correctly from their point of view (which approximately equals p_mM/N), plus the proportion of Elite voters who vote incorrectly from their point of view (which approximates to $(1-p_e)E/N$). The Mass position $\mathcal{M}$ is expected to defeat the Elite position $\mathcal{E}$, therefore, if and only if

$$p_mM/N + (1-p_e)E/N > p_eE/N + (1-p_m)M/N \qquad \text{Eq. 1}$$

or

$$p_m > (E/M)(p_e - ½) + ½ \qquad \text{Eq. 2}$$

The upshot of Equation 2 is that, even if they are less competent (defined as mistaking their own true interests more often), the Masses can nonetheless prevail over more competent Elites by virtue of their greater numbers. Suppose, for example, the Elites are one-fifth as numerous as the Masses in a large population, and suppose that each member of the Elites is on average p_e=0.70 likely to vote in his own true interests. The position in the

true interest of the Masses, $\mathcal{M}$, is more likely than not to win just so long as each member of the Masses is p_m>0.54 likely to vote for that position himself.

From Equation 2 we know what happens in the limiting case, where the number of voters approaches infinity. While we are certainly very interested in what happens in very large group settings like that, we are also interested in what happens in the context of smaller (e.g., factory-sized[6]) groups. So next let us estimate that.

Table 1 tells us how likely majorities for $\mathcal{M}$ are for given group sizes and different levels of Elite and Mass competence. (Cells in which the inequality from equation 2 holds are marked in the table with an asterisk.) Table 1 confirms that what is true for large numbers also tends to be true for smaller numbers: the position in the interest of the Masses is more likely to prevail where the Masses are substantially more numerous or not much less competent than Elites – but not otherwise.

[Table 1 about here]

Take the case from Table 1 that is the most analogous to the one just discussed, where the Elite has size E=21 and competence p_e=0.7 and the Mass has size M=100. From Table 1 we see that $\mathcal{M}$ (the position in the true interest

[6] Assume the factory is a cooperative, so decisions are made by a vote among all the members working in that factory. But assume some members work in the management and others work on the shop floor, and each of those groups have differing interests.

of the Masses) is 56% likely to win if the Mass competence is p_m=0.55, but it is only 28% likely to win if p_m=0.51. Despite members of the Masses still being individually more likely to be right than wrong, they are not so by a sufficiently wide margin in that latter case for the Masses to prevail by sheer weight of numbers alone.

In short: Despite their lower individual competence, the Masses can sometimes win by force of the sheer weight of numbers alone, without any coordination whatsoever. But that happens only within strict limits. Overcoming those limits is where the strategy of epistemic solidarity comes into play.

II. Introducing Epistemic Solidarity

A. Masses against Elites

Suppose that all members of the Masses can recognize one another perfectly. Suppose that all members of the Masses make a pact, to which they all adhere perfectly, to vote the same way in the election. Suppose that they determine which way that will be by a majority vote in a pre-election ballot among the Masses. In that pre-election ballot, every member of the Masses votes sincerely and independently of every other, just as before. But come the

subsequent election itself, all members of the Masses vote, by that institutional arrangement, as a completely unified bloc. That is how we envisage the strategy of 'epistemic solidarity' working, in practice.

Ex hypothesi, each member of the Masses votes in the pre-election ballot independently of each other. Ex hypothesi, members of the Masses are more likely than random to be correct about where the interests of the Masses (which is the same for every member of the Masses) truly lie. Ex hypothesi, there are a great many members of the Masses. So the conditions of the CJT will obtain, and we can be broadly confident that the majority vote in the pre-election ballot among the Masses indicates where the true interests of the Masses lie, just so long as Mass voters vote sincerely in line with their private signals in that ballot.

The literature on strategic voting tells us that sincere voting is not necessarily (or even typically) a Nash equilibrium. Then again, universal strategic voting is often not a Nash equilibrium either.[7] Typically the Nash equilibria that do exist are not easily understood or anticipated, and hence not very likely to emerge among boundedly rational actors who have limited time, attention, information and cognitive capacities.

But even if strategic Nash equilibria are practically unlikely to emerge, one may wonder whether it is plausible that out-of-equilibrium sincere voting

[7] In the classic Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1998) set-up, if everyone votes strategically then nothing can be learned from the assumption that one's vote is pivotal – in the worst case, no one takes into account any private information.

is behaviorally stable. We argue that it will be. The voting game being played will typically be such that sincere voting is, in fact, a Nash equilibrium because of the combination of two facts: insincerity is punished; and the incentives for strategic voting are small.

To see the first, note that the very idea of epistemic solidarity presupposes truthful revelation of one's private signal in the pre-election ballot. Given the purpose for which the Masses institute the pre-election ballot, a strong norm of sincere, non-strategic voting in the pre-ballot is likely to emerge; and those who are seen to deviate from that norm are likely to suffer social sanctions, if only reputational damage. Even if votes are secret and unobservable, the costs in terms of negative self-image or the costs of pretending to have voted sincerely can tip the balance. To see the second fact, note that the chances of any particular vote being decisive among a large group are small, and the incentives for strategic voting based on pivotality considerations are therefore limited. It is exceedingly unlikely you will be the pivotal voter in large populations, which results in a small expected gain from strategic voting. By contrast, the threat of sanctions to ensure truth-telling in the pre-election ballot can be powerful, and is likely to outweigh the strategic incentives. Taking these factors into account, it is plausible that sincere voting is indeed a Nash equilibrium.

In the limiting case (where the size of the Masses approaches infinity), we can be completely confident that the pre-election ballot among the Masses

will indicate where the true interests of the Masses lie, assuming no (or at least not too much) strategic voting occurs in it and the other CJT assumptions are met.[8] Since the probability that the majority is correct increases rapidly with the number of voters for any competence level much above ½, this result approximately obtains even among much smaller groups. Table 2 displays the probability that the majority vote among groups numbering between 40 and 100 members will be correct, for varying levels of individual competence. There we see, for example, that even for a group numbering merely 100, if the individual competence of members of that group is p_m=0.55 the probability that a majority among them is correct is 0.841.

[Table 2 about here]

Suppose now that all members of the Masses practice epistemic solidarity, by voting in the election for whatever option won the pre-election ballot among their group. Then, as long as the Masses' pre-election ballot succeeds in correctly picking $\mathscr{M}$ (which we have just seen is very likely, even among relatively modest-sized groups), the Masses' preferred alternative is highly likely to win in the subsequent election. At the limit, with population

[8] With the important caveat, of course, that the voters are independent, conditional on the state. Neither the convergence to 1 nor the monotonic increase of group competence in group size necessarily obtains once the independence condition is weakened. See Dietrich and Spiekermann (2013) for a discussion and a theorem in that regard.

size going towards infinity while keeping the ratio E:M fixed, the proportion of votes for $\mathcal{M}$ (the position in the true interest of the Masses) will be $[M+(1-p_e)E]/N$ and the proportion of votes for $\mathcal{E}$ will be p_eE/N. Since, ex hypothesi, $M>E$, the position of the Masses would prevail comfortably.

Thus, the practice of epistemic solidarity can be a powerful tool in the hands of the Masses. But in one way, it might look like a tool of strictly limited utility. Not only does its use presuppose that both competence and independence assumptions hold. Furthermore, it is only within a fairly narrow range of values of p_m that the tool will at one and the same time both work and be needed. It will work only where $p_m>0.5$ (with competence below 0.5 the theorem's optimistic conclusions no longer follow). And it is needed only where the Masses would not win by the force of numbers alone, which at the limit is where $p_m > (E/M)(p_e - ½) + ½$. In the example sketched above (where $p_e=0.7$ and $E/M=0.20$) that value would be $p_m < 0.54$. In that example, therefore, the Masses both need and stand to benefit from epistemic solidarity only within a relatively narrow range $0.50<p_m<0.54$. Still, many real world cases may well fall within that window – which is to say, the Masses might often be better than random, but not by much.

Furthermore, it is worth noting the magnitude of the contribution that epistemic solidarity can make among smaller populations to the probability of a victory for the position in the interests of the Masses. Table 3 pulls together for ease of comparison values from table 1 and 2. It displays in the

bottom right corner of each cell the probability of a victory for $\mathcal{M}$ (the position in the true interest of the Masses) for the case of E=21 and M=100, for various values of p_m, assuming the Masses practice epistemic solidarity and the Elites do not. For ease of comparison, the probability of a victory for $\mathcal{M}$ if neither Masses nor Elites practice epistemic solidarity is reproduced as the italicized number in the upper left corner of each cell.

[Table 3 about here]

First concentrate on the column of Table 3 where p_m=0.51. For all the values reported, practicing epistemic solidarity will probably make literally the difference between the Masses winning and losing. We knew that much from Equation 2. But the thing to notice from Table 3 is how very much of a difference it makes to the probabilities, even in this relatively small-group setting. Take once again the case of p_m=0.51 and p_e=0.7. Without epistemic solidarity the chance of an $\mathcal{M}$ victory is only 28%; with it, that likelihood jumps to 58%. Furthermore, practicing epistemic solidarity makes more of a difference the larger the individual competence gap between Elites and Masses. Take the case of p_m=0.51 and p_e=0.8. Without epistemic solidarity $\mathcal{M}$ has almost no chance of winning (the likelihood of that is only 16%), whereas with the Masses practicing epistemic solidarity that likelihood jumps to 58% again.

Next take a case where epistemic solidarity is not strictly needed, in the sense that the position most in the interest of the Masses is likely to win anyway. Consider for example the cell in Table 3 where p_m=0.55 and p_e=0.70. Even if the Masses do not practice epistemic solidarity, $\mathcal{M}$ is 56% likely to win. But if they do, that likelihood jumps to 84%. Politically, that is a huge difference – the difference between a close-run thing and a virtual certainty. So even in these sorts of cases, the Masses can benefit greatly from practicing epistemic solidarity, even in relatively small-group settings.

B. Elites Against Masses

Of course, either side or both could avail themselves of the strategy of epistemic solidarity. Conventionally, solidarity is most often discussed as a weapon of the weak, not least because they are in most need of it to overcome the strong. But solidarity may actually be practiced more easily among the strong, who are better networked and who thus find it easier to exchange information and coordinate their actions to ensure their interests are served. For these practical reasons, epistemic solidarity (as opposed to other forms of solidarity, perhaps) may turn out to be a weapon more available to the Elites.

Suppose both the Elites and the Masses practice epistemic solidarity within their own groups. Then the law of large numbers tells us that in the

limiting case (as both E and M approach infinity, keeping their ratio constant) the proportion of votes for each position would simply equal the proportion of members of each group. With a proportion of approximately M/N votes for $\mathcal{M}$ and E/N votes for $\mathcal{E}$, and $M>E$, the Mass position would ordinarily be the clear winner. That follows straightforwardly, and is relatively uninteresting.

More interesting is the case in which the Elites practice epistemic solidarity while the Masses do not. That compounds the epistemic advantage that the more competent Elites already have over the Masses. Among large populations (with fixed proportion E:M), vote proportions would be approximately $(E+[1-p_m]M)/N$ for $\mathcal{E}$ and p_mM/N for $\mathcal{M}$. Thus, the Elites practicing epistemic solidarity increases the number of votes for $\mathcal{E}$ and reduces the number for $\mathcal{M}$, compared to the case where neither group practices epistemic solidarity. At the limit, the Elites win if

$$p_m < E/2M + ½ \qquad \text{Eq. 3}$$

For example, suppose, as before, M is five times larger than E, and suppose the Elites practice solidarity and the Masses do not. Then at the limit the option that is in the interests of the Elites, $\mathcal{E}$, will prevail whenever $p_m < 0.6$.

Table 4 provides a few examples for smaller populations with the same group size and competence parameters as in Table 1. As we see from Table 4,

even moderately small Elite groups practicing epistemic solidarity can seriously reduce the probability of a win for the option that is in the Masses' interests. Take the case discussed above, in which the Elite has size E=21 and individual competence p_e=0.7 and the Mass has size M=100 and individual competence p_m=0.55. From Table 4 we see that, if Elites practice epistemic solidarity while the Masses do not, the probability of $\mathcal{M}$ (the position in the true interest of the Masses) winning is only 16%. That compares to 56% probability of $\mathcal{M}$ winning when neither Elites nor Masses were practicing epistemic solidarity, as reported in Table 1.

[Table 4 about here]

But even if the Elites practice epistemic solidarity and the Masses do not, that does not always lead to an Elite victory. Look what happens, for example, if the Elites are much smaller relative to the Masses. Consider the case of E=11 and M=100, with the same levels of individual competence as before. Then there would be a 51% chance of $\mathcal{M}$ (the position in the true interest of the Masses) winning, despite the Elites practicing epistemic solidarity and the Masses not.

III. Interim Conclusions

Despite being substantially more numerous than the Elites, the Masses might nonetheless lose to them because individual members the Masses are substantially less competent at judging their true interests. But as we have shown, the Masses can often rectify that by practicing epistemic solidarity, pooling their information about their interests with one another. If they do so they will typically prevail over the Elites, whether or not the Elites do the same. But if the Elites practice epistemic solidarity while the Masses do not, the Elites can sometimes in that way beat the Masses.

Let us summarize these results with the aid of a numerical example. Imagine a society composed of an Elite numbering 200,000 and 1 million members of the Masses. Suppose that the competence of individual members of the Elite in judging their own true interests is p_e=0.7 while that of individual members of the Masses is p_m=0.51. The approximate number of votes that can be statistically expected for each option from each type of voter is as shown in Figure 1 as the gray and white proportion of the bars.

[Figure 1 about here]

In that Figure 1 example, the Masses lose narrowly if neither they nor the Elites practice epistemic solidarity, and the Masses lose by an even wider margin if the Elites practice solidarity and the Masses do not. The figure also

shows why: if the Masses do not pool their information, 49% of them (the hatched white bar section) mistakenly support $\mathcal{E}$ instead of $\mathcal{M}$. But just so long as the Masses themselves practice epistemic solidarity, the Masses prevail. And that remains almost as powerfully true whether or not the Elites practice epistemic solidarity as well.

IV. Sensitivity to Uncertainty Concerning Who Belongs in the Group

As we have seen, people who have the same interests but are individually not terribly competent in identifying what serves their interest (like the Masses) can find out with great reliability what is in their interest if they take a majority vote among themselves. In that way, the Masses can usually succeed in outvoting the Elites. However, in order to do that, they first have to identify 'who is with them'.

That may well be a challenge for the Masses in particular. If they are individually not very good at identifying what is in their interest, they may also find it difficult to know with whom they share the same interests. The Elites, by contrast, may have a few aces up their sleeves: in addition to being more competent individually, they might be able to devote more efforts to finding out who is 'with them', they tend to 'know people who know', they are probably socially more mobile and better networked, and they often dominate the public discourse. All this helps the Elites to identify their own

and to vote for their interests as a block. As we have seen, if the Masses remain divided while the Elites coordinate their votes, the Elites will often be able to impose their minority interests on the community as a whole.

So far we have been assuming that people have perfect information about who is in the group that shares the same interests as they do. If so, then the group with which they practice epistemic solidarity will contain all and only those with whom they share an interest. In the real world, however, there is bound to be some uncertainty surrounding who shares the same interests with them. Just how sensitive might our findings be to those uncertainties?

A. Incomplete Assortation: Some Abstain from Epistemic Solidarity

Basically, there are two different ways an agent might respond to uncertainty over which group shares his interests. Someone who is subjectively particularly uncertain and averse to the risk of joining the wrong group might prefer to abstain from practicing epistemic solidarity with either group. Abstaining means not joining a group, not taking part in a pre-ballot and voting purely on the basis of one's own individual judgment of where one's own interests lie.[9]

[9] Another theoretical possibility is to join a group and take part in their pre-ballot, but then vote according to one's private signal. Note, however, that we assumed that once one has joined a pooling group voting in line with the pre-ballot is institutionally

If some individuals abstain from epistemic pooling, this would simply create a situation in between that represented by the values in the two corners of each cell in Table 3. The top left value in each cell there represents the probability of an $\mathcal{M}$ victory if none of the Mass voters practiced epistemic solidarity; the bottom right value represents the probability of an $\mathcal{M}$ victory if all the Mass voters practiced epistemic solidarity. If, for instance, only half of the Mass voters practiced epistemic solidarity, the probability of an $\mathcal{M}$ victory would be in between those two values (tilted towards the higher value, as the marginal returns of pooling are decreasing in group size). In the case that served as our previous running example of M=100, E=21, p_e=0.7 and p_m=0.55, the probability of an $\mathcal{M}$ victory if only half the Masses practice epistemic solidarity is 0.76. Thus, it is not always necessary for *all* of the Masses to pool their information to win.

For very large populations, the outcome at the limit can be calculated in the same manner as before. Suppose a fixed proportion of φ_M of the M members of the Masses practice epistemic solidarity (with $1>\varphi_M>0$) and the rest of the Masses vote on the basis of their own individual judgment. As before, we assume that $p_m>0.5$, so that each Mass voter has the same better-than-random probability of individually correctly assessing where his true

---

required. Influencing the pre-ballot without following it is therefore not possible. This is a plausible restriction, as pooling groups would be likely to have strong norms (which people joining the group themselves internalize) against members who enter their vote in the pre-ballot without following it. We have invoked precisely such a norm in our argument against strategic voting above.

interests lie. Similarly, each member of the Elites has competence p_e>0.5 of individually correctly assessing where her true Elite interests lie, and suppose none of the Elites practice epistemic solidarity. Then at the limit, as population size goes to infinity while keeping the ratio E:M fixed, the pre-election ballot will direct a share of approximately $\varphi_M M/N$ voters toward $\mathscr{M}$; and $\mathscr{M}$ will garner approximately another $p_m(1 - \varphi_M)M/N$ proportion of the votes from members of the Masses not practicing epistemic solidarity and approximately a $E(1 - p_e)/N$ proportion of votes from members of the Elites mistakenly voting against their own true interests. $\mathscr{E}$ will garner approximately a $p_e E/N$ share of votes from members of the Elites voting correctly in their true interests and approximately another $(1-p_m)(1 - \varphi_M)M/N$ share of votes from members of the Masses who do not practice epistemic solidarity, voting mistakenly against their own true interests. Thus, at the limit, $\mathscr{M}$ will defeat $\mathscr{E}$ if

$$\varphi_M M + p_m(1 - \varphi_M)M + E(1 - p_e) > p_e E + (1 - p_m)(1 - \varphi_M)M , \qquad \text{Eq. 3}$$

which can be rearranged to

$$\varphi_M > [(E/M)(p_e - \tfrac{1}{2}) - (p_m - \tfrac{1}{2})]/[(1 - p_m)]. \qquad \text{Eq. 4}$$

Thus, for example, if p_m=0.51 and p_e =0.7 and E and M are both large with E:M=1:5, then $\mathcal{M}$ is expected to win so long as a little over 6% of the Masses practice epistemic solidarity and none of the Elites do.

Next suppose that both the Elites and the Masses practice epistemic solidarity, but some of each abstain from that practice on grounds they are subjectively too uncertain which is their own true group. Suppose once again that $\varphi_M M$ out of the total M true members of the Masses practice epistemic solidarity; and now add to that the assumption that $\varphi_E E$ out of the total E true members of the Elites practice epistemic solidarity (with $1>\varphi_E>0$). Those not practicing epistemic solidarity vote on the basis of their individual perception of where their true interests lie, with accuracy of p_m and p_e for members of the Masses and Elites, as before.

By reasoning analogous to that underlying Equation 3, at the limit $\mathcal{M}$ is expected to beat $\mathcal{E}$ if

$$\varphi_M M + p_m(1 - \varphi_M)M + (1 - p_e)(1 - \varphi_E)E$$
$$> \varphi_E E + p_e(1 - \varphi_E)E + (1 - p_m)(1 - \varphi_M)M, \qquad \text{Eq. 5}$$

which can be rearranged to

$$\varphi_M > [(E/M)(\varphi_E + p_e - p_e\varphi_E - ½) - (p_m - ½)] / (1 - p_m). \qquad \text{Eq. 6}$$

That means that, in a similar scenario to the one just considered (p_m=0.51, p_e =0.7, E:M=1:5) then if just half of the true members of the Elites practice

epistemic solidarity, $\mathcal{M}$ is expected to win so long as more than about 12.2% of the Masses practice epistemic solidarity. Even if 80% of the Elites practice epistemic solidarity, only more than about 15.9% of the Masses need to do so in order to make an $\mathcal{M}$ victory more likely than not.

Inequalities 4 and 6 come in handy if we want to explore how sensitive our conclusions are to abstentions from epistemic solidarity. In our running example, as long as a non-negligible proportion of the Masses practice epistemic solidarity, the option in the Masses' true interests will win, and that is true within broad limits no matter how many of the Elites practice epistemic solidarity. This also becomes clear by looking at the large grey hatched bar when the Masses pool in Figure 1: their pooled votes carry the Masses comfortably over the majority threshold, so that there is a lot of room for less pooling discipline without a change in outcome. However, different parameter values might put the result much more on a knife's edge, so that near universal pooling would be required.

B. Imperfect Assortation

A second possible response, tempting to those who are subjectively uncertain but perhaps not quite so uncertain or not quite so risk averse, is to practice epistemic solidarity with the group that they think is most likely to share their own interests – knowing that there is a risk they will get that assessment

wrong, and end up practicing epistemic solidarity with the 'wrong' group, from their own point of view.

The groups in which pooling takes place would then no longer be homogeneous, as they were (by stipulation) in the models discussed previously.

For the purpose of this model, assume that everyone knows that there are exactly two types of people in the population. One is the Mass type, the other is the Elite type, and just as before there are M of the former and E of the latter. Let us further assume that all Elite type individuals have the same probability $pg_e > 0.5$ of correctly identifying which type they are, while all Mass type individuals have probability pg_m. Call this the 'group selection competence' of the Mass and Elite type, respectively. Let the population then be exhaustively partitioned into two groups, one composed of self-assessed members of the Masses and the other self-assessed members of the Elites. Note that the sizes and compositions of these groups can vary, as they are the result of a stochastic assortation process.

Logically, there could be strategic considerations standing against the truthful revelation of one's perception of one's group type.[10] But here we rule

[10] A notable Nash equilibrium has *all* individuals end up in the same group with a pooling pre-ballot. Any unilateral deviation is unattractive, as the large pooling group always wins against one voter in the other group, while being in the large groups provides a non-zero probability of being pivotal. In fact, if the larger group is a pooling group, being in the smaller group is dominated by being in the larger, winning group. This may be of some real-world interest: if individuals expect that one group will be larger *and* a pooling group, then this group is preferred if the individuals only care about getting their preferred result.

out strategic considerations, in terms of group choice as well as pre-ballot voting. This is not purely for convenience of modeling. There may be good sociological reasons for people to reveal truthfully their perception of to which group they belong. They may have an expressive desire to join 'their own group' or, as before, a normative commitment to positively contribute to epistemic pooling within their own true group. In addition, joining the 'right' group may have positive side-payments that outweigh any strategic considerations.

Finally, suppose that all Mass (respectively: Elite) type individuals have probability p_m>0.5 (p_e>0.5) of being correct in their personal assessment of their own interests in the case at hand, as before. In the first instance, we can explore this setup with computer simulations, investigating how the group selection competence influences the epistemic success of the Elites and Masses.

In Figure 2, we plot the proportions of Mass majorities (relying on 1000 simulations for each data point) as a function of group selection competence, which for now we assume to be equal for both types, such that $pg_m = pg_e$. The number of Elite types is 21 and of Mass types 100. The former have competence $p_e = 0.7$ and the latter $p_m = 0.55$. The circle markers show the probability of a Mass majority when only the self-assessed Elite group pools, the diamonds when only the self-assessed Mass group pools, and the stars when both groups pool.

[Figure 2 about here]

We know from Table 4 that if the Elites and the Masses self-identify completely correctly and the Elites alone pool their votes, the probability of a Mass victory is 16%. This result is reflected in Figure 2 by the right-most circle marker: when group selection competence is 1, Mass majorities have a probability of about 16%. It is, prima facie, unsurprising that the Elites benefit from higher group selection competence when they are the only group pooling. By contrast, when the Masses or both groups pool votes, then the larger size of the self-assessed Mass group turns a higher group selection competence into an advantage for the Masses – the more homogeneous the pooling groups become, the more epistemically successful the Masses become in their pooling, outvoting the small Elite group quite reliably.

This looks like a straightforward story. But consider Figure 3, which is the same as Figure 1 except with competence parameters altered to p_e=0.8 and p_m=0.6. There, an interesting twist to that story stands out better. Focus on the curve of circle markers (that is, pooling of the self-assessed Elites only). The probability of a Mass majority is at its lowest at a group selection competence of about 85%. That suggests that, when only the self-assessed Elites pool their votes, the Elites benefits most from individuals making occasional mistakes when choosing their group.

[Figure 3 about here]

The reason lies in the variable sizes of the self-assessed groups. Were group selection competence set to 1, all Elite types would end up in the Elite group and all Mass types in the Mass group, leading to group sizes *E* and *M*. However, if group selection competence is below 1 (but above 0.5) and *E*<*M*, we would expect the self-assessed Elite group to be larger than *E* and the self-assessed Mass group to be smaller than *M* because there will be more Masses who mistakenly choose the Elite group than there are Elites who mistakenly choose the Mass group. This increase in the size of the Elite group benefits the pooling Elites because (as long as the proportion of truly Mass agents in the self-assessed Elite group is small enough to be outvoted reliably by the true Elites) the Elite group in this way 'captures' some unsuspecting Mass voters and, by pooling, leads them to vote for the Elite interests.[11]

So far we have been assuming group selection competence is identical for everyone in the population. Next let us see what happens if we hold that constant for the Masses, at pg_m=0.6, while letting the group selection of the

[11] If we were allowing strategic behavior (which here we are assuming away), that would suggest a strategy for the Masses: if a great many of them could strategically coordinate to pose as members of the Elite and vote in the Elite's pre-election ballot in line with their own true interest $\mathcal{M}$ in sufficient numbers for that to win the Elite's pre-election ballot, they could in that way hijack the Elite's epistemic pooling in the service of their own true Mass interests. Of course as soon as the Elite realized this was happening the rules of their epistemic pooling would probably change, so that e.g. members of the pooling group cannot simply self-nominate as members but instead would have to be accepted by sufficiently many other members.

Elites pg_e vary. The results of that are displayed in Figure 4, for the case once again of M=100 E=21, p_e=0.7 and p_m=0.55.

[Figure 4 about here]

Two things change between Figures 3 and 4. First, the rise at the end of the row of circle markers (where only the Elites pool) disappears. That is just as we would expect, given our explanation for the rise that was observed in Figure 3. That, we argued, resulted from fewer Mass agents mistakenly identifying themselves as Elite as the group identification competence of the Masses (as well as of the Elites) increases in Figure 3. But in Figure 4, pg_m is held constant at 0.6, so roughly the same proportion of Mass agents will mistakenly join the Elite group across all cases shown in Figure 3.

Second and more interesting is what happens in the row of star markers (where both Elites and Masses pool). If the Elites are more competent at recognizing their true type, then even where *both* Elites and Masses practice epistemic solidarity the Elites benefit more from that practice.

Indeed, very high group selection competence among the Elites might even lead to an Elite victory, despite the fact that Elites and Masses are both pooling. With the parameters set as in Figure 4, for example, the row of star markers gets close to the 0.5 threshold for values of pg_e around 0.95. That is,

however, obviously a very extreme case, involving the unrealistically high value of pg_e~0.95.

The upshot of our analysis in this section is that our interim conclusions stand up well and can be robust to the introduction of uncertainty regarding group choice. There are basically two types of responses to such uncertainty. One is to abstain from practicing epistemic solidarity at all. The other is to take one's chances, practicing epistemic solidarity with whichever group seems most likely to be truly your own but knowing you might be wrong about that. Our analysis suggest that, depending on the parameters, a very substantial proportion of the Masses can abstain in the first way, or be more likely to get it wrong in the second way than the Elites, and our overall conclusion still stands up.

V. Conclusion

Solidarity is often taken as a matter of concerted action: pooling resources or coordinating behavior. The Masses improve their chances of overcoming the smaller but more powerful Elites if they display solidarity in that sense. We have shown that solidarity can also be about pooling in quite a different sense: the joint formation of correct beliefs. The Masses may be uncertain about what is truly in their interest, and if they succeed in pooling the

dispersed pieces of information they hold they can overcome this 'false consciousness'. This strategy can work well, but it faces an obvious problem: to successfully identify the Mass interest by information pooling, the Masses need to know who they have a shared interest with. If they fail to identify their own, while the Elites succeed, the well-organized Elites may gain the upper hand, even though they are much smaller in numbers. Our results give a new twist to the old adage that 'knowledge is power' – one needs to know one's own interest, but to acquire that knowledge, one needs to know who knows.

References

Austen-Smith, David and Jeffrey S. Banks. 1996. Information aggregation, rationality, and the Condorcet Jury Theorem. *American Political Science Review* 90 (1): 34-45.

Dietrich, Franz and Kai Spiekermann. 2013. Independent opinions? On the causal foundations of belief formation and Jury Theorems. *Mind* 122 (487): 655-685.

Goldman, Alvin. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.

Goodin, R. E. and K. W. S. Roberts. 1975. The ethical voter. *American Political Science Review* 69 (3): 926-28.

Feddersen, Timothy and Wolfgang Pesendorfer. 1998. Convicting the innocent: the inferiority of unanimous jury verdicts under strategic voting. *American Political Science Review* 92 (1): 23-35.

Grofman, Bernard; Guillermo Owen; and Scott L. Feld. 1983. Thirteen theorems in search of the truth. *Theory & Decision* 15: 261-78.

List, Christian and Robert E. Goodin. 2001. Epistemic democracy: generalizing the Condorcet Jury Theorem. *Journal of Political Philosophy* 9: 276-306.

	p_m=0.51				p_m=0.55				p_m=0.6			
p_e=0.6	E/M	40	60	80	100	E/M	40	60	80	100	E/M	40	60	80	100
	11	0.42	0.45	0.47	0.49	11	0.60*	0.68*	0.73*	0.77*	11	0.80*	0.88*	0.93*	0.96*
	21	0.33	0.37	0.40	0.42	21	0.49	0.58*	0.65*	0.70*	21	0.69*	0.81*	0.88*	0.93*
	31	0.26	0.30	0.33	0.36	31	0.40	0.49	0.57*	0.63*	31	0.59*	0.73*	0.83*	0.89*
p_e=0.7	E/M	40	60	80	100	E/M	40	60	80	100	E/M	40	60	80	100
	11	0.30	0.35	0.38	0.41	11	0.48	0.58*	0.65*	0.71*	11	0.70*	0.82*	0.89*	0.94*
	21	0.16	0.21	0.25	0.28	21	0.28	0.39	0.48	0.56*	21	0.48	0.66*	0.78*	0.86*
	31	0.08	0.11	0.15	0.18	31	0.15	0.24	0.33	0.42	31	0.29	0.48	0.64*	0.75*
p_e=0.8	E/M	40	60	80	100	E/M	40	60	80	100	E/M	40	60	80	100
	11	0.20	0.25	0.30	0.33	11	0.35	0.47	0.56*	0.63*	11	0.58*	0.75*	0.85*	0.91*
	21	0.05	0.09	0.13	0.16	21	0.12	0.22	0.32	0.40	21	0.26	0.47	0.64*	0.76*
	31	0.01	0.03	0.04	0.06	31	0.03	0.08	0.14	0.22	31	0.08	0.23	0.40	0.55*

Table 1: Probabilities of a majority for the alternative in the interests of the Masses for different Elite and Mass competence and group sizes. Values with asterisk are those for which inequality (2) is true.

	40	**60**	**80**	**100**
p_m **= 0.51**	0.550	0.561	0.571	0.579
p_m **= 0.55**	0.736	0.780	0.814	0.841
p_m **= 0.6**	0.898	0.940	0.964	0.978

Table 2: Probabilities of majorities for the alternative in the interests of the Masses in the pre-election ballot among the Masses, according to the Condorcet Jury Theorem (assuming that ties are broken by a coin toss).

	p_m = 0.51	p_m = 0.55	p_m = 0.6
p_e = 0.6	*0.42* 0.58	*0.70* 0.84	*0.93* 0.98
p_e = 0.7	*0.28* 0.58	*0.56* 0.84	*0.86* 0.98
p_e = 0.8	*0.16* 0.58	*0.40* 0.84	*0.76* 0.98

Table 3: Probabilities of a majorities for the alternative in the interests of the Masses in election assuming all Mass voters vote strictly in accordance with pre-election ballot among the Masses, assuming *E*=21 and *M*=100.
(Probability without pre-election ballot in top left of cell it italics, probability following pre-election ballot in bottom right of cell.)

	p_m=0.51				p_m=0.55				p_m=0.6			
p_e=0.6	E/M	40	60	80	100	E/M	40	60	80	100	E/M	40	60	80	100
	11	0.28	0.31	0.34	0.37	11	0.35	0.44	0.52	0.59	11	0.48	0.67	0.78	0.86
	21	0.17	0.18	0.18	0.20	21	0.18	0.19	0.23	0.29	21	0.20	0.27	0.41	0.56
	31	0.13	0.13	0.13	0.13	31	0.13	0.13	0.13	0.15	31	0.13	0.13	0.17	0.25
p_e=0.7	E/M	40	60	80	100	E/M	40	60	80	100	E/M	40	60	80	100
	11	0.13	0.16	0.21	0.24	11	0.20	0.31	0.42	0.51	11	0.37	0.59	0.74	0.84
	21	0.03	0.03	0.04	0.05	21	0.03	0.05	0.10	0.16	21	0.04	0.14	0.30	0.47
	31	0.01	0.01	0.01	0.01	31	0.01	0.01	0.01	0.03	31	0.01	0.01	0.05	0.13
p_e=0.8	E/M	40	60	80	100	E/M	40	60	80	100	E/M	40	60	80	100
	11	0.07	0.11	0.15	0.19	11	0.14	0.27	0.38	0.46	11	0.33	0.57	0.73	0.82
	21	0.00	0.01	0.02	0.03	21	0.00	0.03	0.08	0.13	21	0.02	0.12	0.29	0.46
	31	0.00	0.00	0.00	0.00	31	0.00	0.00	0.00	0.02	31	0.00	0.01	0.04	0.13

Table 4: Probabilities of a majority for the alternative in the interests of the Masses for different Elite and Mass competence and group sizes, with Elites pooling their votes (based on 10,000 vote simulations each).

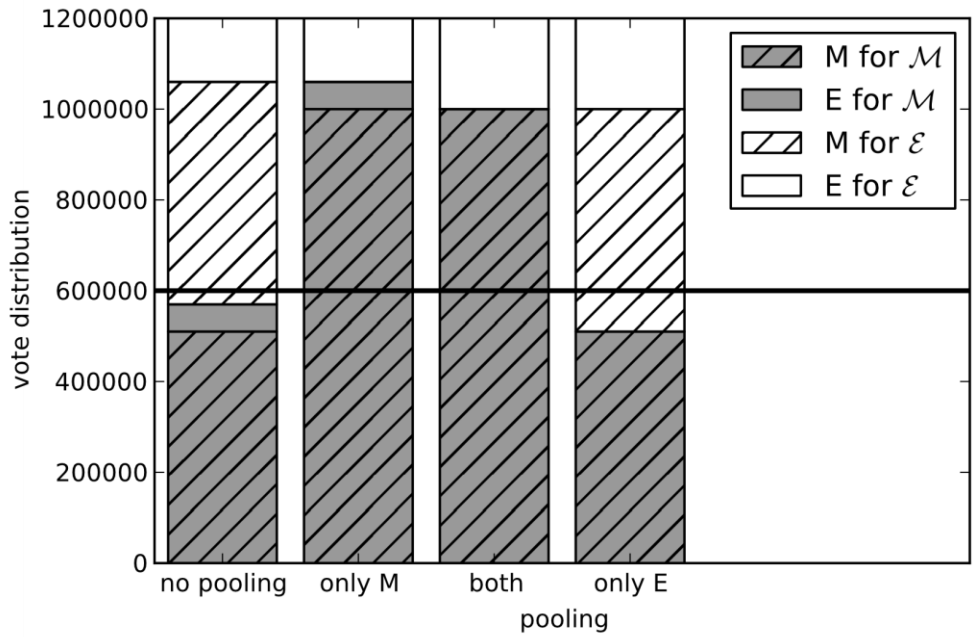

Figure 1: Approximate expected vote distribution, E=200,000, M=1,000,000; p_e=0.7, p_m=0.51.

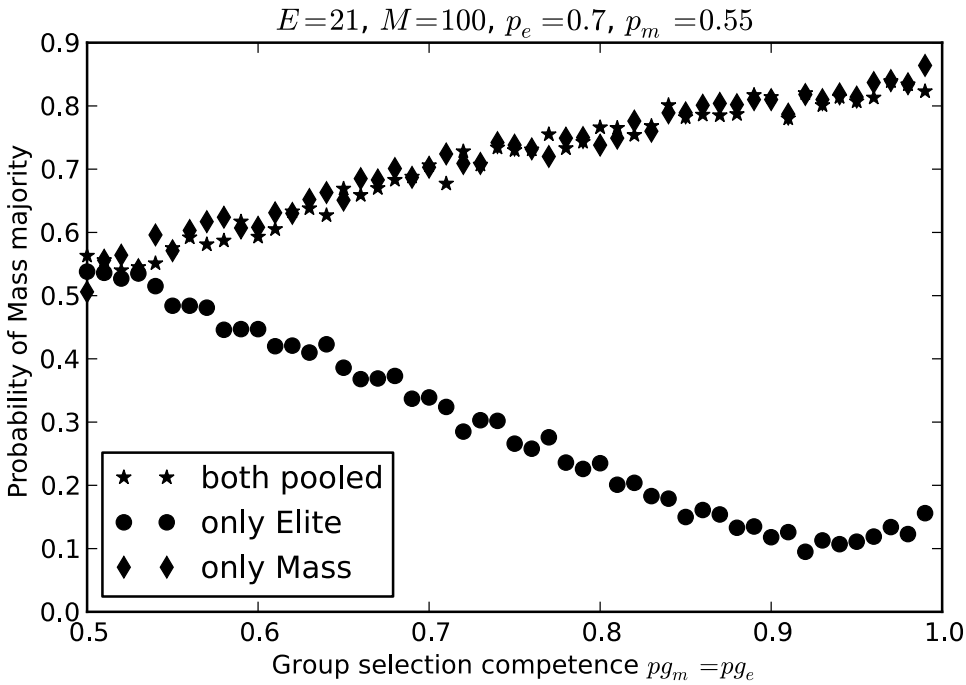

Figure 2: Probability of Mass majorities as a function of group selection competence.

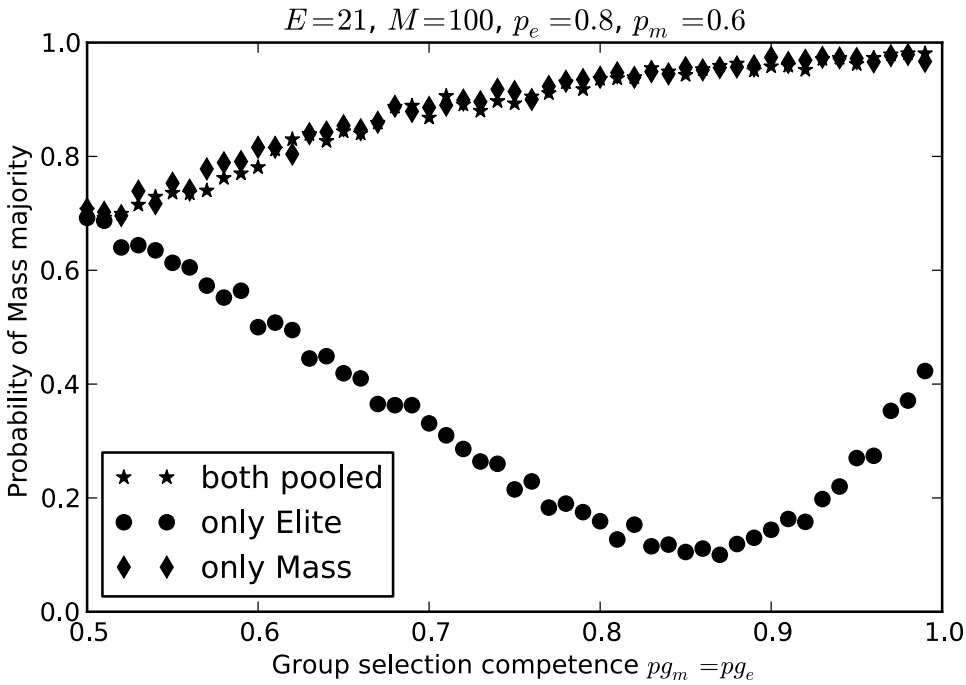

Figure 3: Probability of Mass majorities as a function of group selection competence.

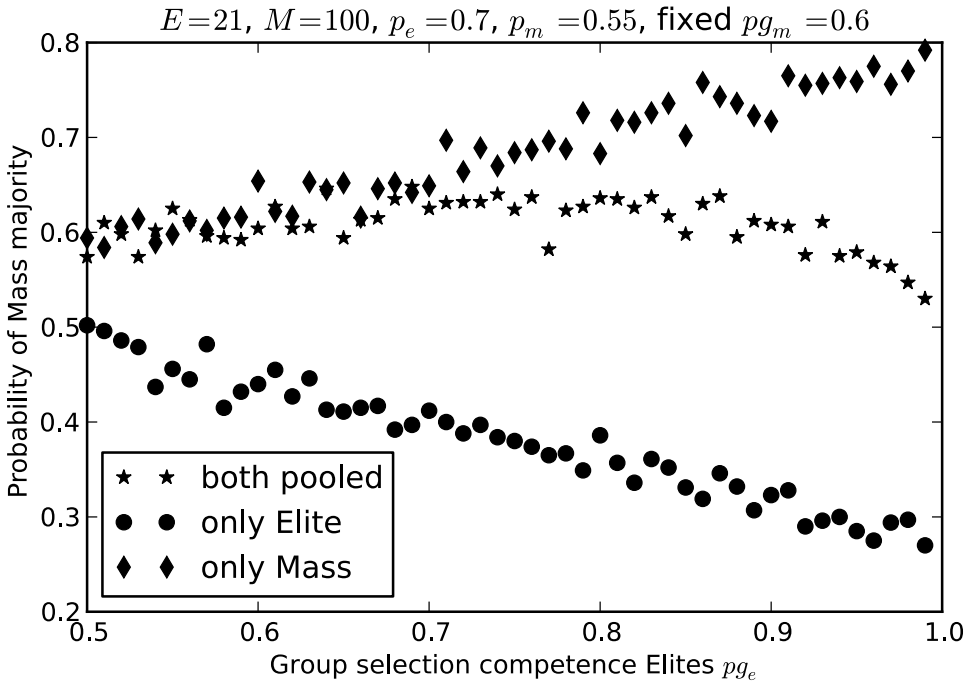

Figure 4: Probability of Mass majorities as a function of the Elite group selection competence, Mass group selection competence fixed.