The Value of Formalism: Re-examining External Costs and Decision Costs with Multiple Groups

Keith L. Dougherty
Department of Political Science
University of Georgia
dougherk@uga.edu

Julian Edward
Department of Mathematics
Florida International University
edwardj@fiu.edu

Robi Ragan
Department of Economics
San Jose State University
robi.ragan@sjsu.edu

August 5, 2013

Abstract

Several authors have examined the optimal k-majority rule using a variety of criteria. We formalize and extend the original argument laid out by Buchanan and Tullock (1962) using a decision theoretic analysis from the perspective of an individual voter. In our model voters are members of one or more groups, representing different religious or ethnic backgrounds, or clusters of people with similar preferences. We consider both up or down votes on a single proposal as well as voters over a series of proposals. We find that the optimal k-majority rule depends on a number of parameters, most notably the number of rounds needed to create a proposal that will pass under each k-majority rule. In some cases, the external cost function can actually increase over a range. In others, larger k-majority rules can fail to pass Pareto preferred proposals. Our results should help advance the discussion of a classic work in Virginia Political Economy.

1 Introduction

One of the first, and most fundamental, questions for those writing a constitution is to determine how many individuals have to agree in order for the constitution to be enacted. If the threshold is too large, then little progress will be made. If the threshold is too small, then a majority of individuals could impose their will on others. In their classic work, *The Calculus of Consent* (1962), Buchanan and Tullock framed the question as a study of the optimal k-majority rule, which is loosely the threshold of "yea" votes needed for a proposal to pass. K-majority rules range from the affirmative vote of one individual (rule of one) to the affirmative vote of all individuals (unanimity rule), with majority rule and all the supermajority rules in between. Buchanan and Tullock argued that the optimal k-majority rule is the one which minimizes the sum of external costs and decision costs. If decision costs were negligible, then unanimity rule would be the optimal, according to their argument, because it alone minimizes external costs. If decision costs and external costs were equally important, then the optimal k-majority rule would be some k-majority near majority rule.

Using probabilistic models, Dougherty and Edward (2004, 2011), find that the behavior of external costs and decision costs can differ from Buchanan and Tullock's descriptions. The difference is largely because external costs and decision costs are functions of the probability of passing proposals in a probabilistic model. That probability does not decrease linearly as the k-majority rule threshold increases. Instead, the probability of passing proposals follows a logit-type shape which makes proposals almost certain to succeed over a range of small k-majority rules and almost certain to fail over a range of large k-majority rules, with a steep decline in between. The logit shape affects the external cost and decision costs functions, making the k-majority rule which minimizes external costs a range of k-majority rules near unanimity rule. With decision costs included, a variety of k-majority rules may be optimal, some of which are near majority rule. The justification for this result differs from the argument made by Buchanan and Tullock.

One limitation of Dougherty and Edward's analysis is that they examine the expected number of individuals who will lose under various k-majority rules, rather than the expected loss to an individual.[1] Someone fully committed to methodological individualism might find the latter more enlightening. Another limitation of Dougherty and Edward's analysis is they assume everyone's preferences are drawn from a single unimodal distribution, rather than from different distributions for different groups in the population. These distinctions are important because their analysis seems to be from the perspective of the collective and because a unimodal distribution may unintentionally limit the possibility of passing proposals against the interest of some group.

This paper addresses these shortcomings by examining the optimal k-majority rule from the perspective of a rational decision maker who is uncertain about whether he will gain or lose from a constitution constructed in his/her community. We consider this decision for cases where the distribution of preferences are symmetric and cases where the distribution of preferences in the community are asymmetric. We also consider a case where some individuals are more likely to gain from a collective decision and others are more likely to lose from it as new proposals are considered.

Our analysis is decision theoretic, which means individuals determine the optimal k-majority rule based on expected gains and losses. They do not adjust their behavior to the k-majority rule assigned, as in a game theoretic model. Game theory would be a welcome extension.

We find several interesting results. First, an individual's relative weighting of external costs and decision costs can make either $k = 1$ optimal or $k = N$ optimal. However, even if an individual puts all of his/her weight on external costs, it is not always the case that $k = N$ will be uniquely optimal, as Buchanan and Tullock suggest. For votes on a single proposal external costs often will be roughly zero for a range of k-majority rules near unanimity rule.

---

[1] More precisely, they examine the expected number of BT-losers where a BT loser is someone who votes for a BT-preferred status quo or a BT-preferred proposal and society chooses the opposite alternative (Dougherty and Edward 2011, 60).

All k-majority rules in this range will be optimal. Similarly, for votes on a single proposal decision costs will be minimized for a range of k-majority rules near $k = 1$ because these k-majority rules are equally likely to pass proposals. If an individual puts all of his/her weight on decision costs, a range of k-majority rules in the neighborhood of $k = 1$ would be optimal.

Second, for votes over a series of proposals, the rate at which successive proposals become more attractive can affect the optimal k-majority rule. If each additional proposal improves utility greatly, then large k-majority rules may be optimal because they will not allow proposals to pass with a large number of losers. If each additional proposal improves utility only slightly, then smaller k-majority rules will be optimal because decision costs will grow faster than external costs.

Third, if the progression of a series of proposals makes some groups better off at the expense of others, then external costs can actually increase over a range of mid-sized k-majority rules. Smaller k-majority rules will pass some of the earliest, less harmful proposals, while mid-sized k-majority rules will not pass a proposal until later in the series when losses to the out-group are larger. We also discover a case where large k-majority rules fail to pass Pareto preferred proposals because a particular group makes systematic errors in judgement. In these cases, larger k-majority rules are more likely to forgo Pareto improvements than smaller k-majority rules because larger k-majority rules are more responsive to the errors made by a few.

Surprisingly, the perspective of an individual, as either a typical member of the voting population or a member of the worse off group, has little affect on the optimal k-majority rule for most of the parameters we examine.[2] External costs can be significantly greater for someone in the worse off group, but external costs almost always tend toward zero at the same point for all groups. Thus, the main difference in group memberships is the magnitude

---

[2] There are notable exceptions, such as when there exists a small dissident group.

of the external costs added to the decision costs, which, when summed, typically has a small influence on the optimal k-majority rule.

2 Previous Literature

In a seminal work, Buchanan and Tullock (1962) argued that the optimal k-majority rule minimizes the sum of external costs and decision costs. External costs are the costs an individual expects to endure as the result of the actions of others (Ibid., p. 62). For example, taxing citizens for street repairs that they do not want creates an external cost for citizens who have to pay taxes for the street repairs. Buchanan and Tullock argued that external costs decrease as the the k-majority threshold increases because members of a decisive coalition (i.e. those who create and pass a proposal) might pass proposals against the interests of those outside their coalition. External costs should be greatest if a single individual were authorized to act for the group ($k = 1$), because the individual might tax everyone equally and spend all the revenue only on roads that he/she travels. External costs should be smallest, typically zero, if group action requires the approval of all members of a group ($k = N$), because individuals would approve of taxes for road repair only if they benefited more than the cost of the tax. If everyone made decisions in their interest, only unanimity rule would guarantee zero external costs for everyone. Because individuals are uncertain about whether they will be a member of the decisive coalition when a k-majority rule has to be adopted, they will consider the *expected* costs imposed under each k-majority rule. Meaning, they have to make a decision about the most appropriate k-majority rule without knowing for certain whether they would gain or lose under the voting rule chosen.[3]

In contrast, decision making costs are the costs resulting from the time and effort needed to reach an agreement. Buchanan and Tullock argue that such costs increase as the k-majority rule threshold increases. This is because very little time and effort is needed for

---

[3] See Heckelman and Dougherty (2010) for a crude test of whether larger k-majority rules have negative effects on tax increases.

one individual to make a decision. More time is required as the k-majority rule threshold increases. In the street repair example, requiring the approval of only one individual to make decisions may lead to quick decisions about street repairs. Requiring the approval of a few more individuals will require a little more time and effort to create a successful plan. If everyone must approve, a considerable amount of time and effort is needed to make everyone satisfied with the proposal.

Buchanan and Tullock illustrate such costs using a figure similar to Figure 1. The number of individuals required to make a decision are represented along the horizontal axis and the expected costs of a particular decision are represented along the vertical axis. The thin solid line, decreasing from left to right, depicts external costs. The thick solid line, increasing from left to right, depicts decision costs. At the left extreme, the rule of $k = 1$ produces potentially large external costs but minimal decision costs. No delays should be expected under such a voting rule. On the right extreme, unanimity rule minimizes external costs but imposes extremely high decision making costs. Buchanan and Tullock suggest that the optimal k-majority rule minimizes the sum of these two functions (depicted by the dashed line). This occurs at $k = 49$ in our figure. Here the sum of the expected costs from a decision being imposed on oneself and the expected time and effort of making a decision are minimized. If legislatures have to chose a voting rule from the set of k-majority rules, $k = 49$ would be optimal in this setting.

Mueller (2003, 76-8) starts with the same basic framework, but adds a "kink" in the decision costs function at $N/2$ which typically makes majority rule the optimal k-majority rule. The kink exists for $k \leq N/2$ because it is possible for both policy x and for policy $\sim x$ to pass. Suppose for example, there are 100 voters and the k-majority rule threshold is set at 40. A proposal to raise taxes for street repairs might receive the support of 43 members and pass. After the measure passes, a counterproposal to eliminate the street repair measure might receive the support of 45 members and also pass. Any k-majority rule, with $k \leq 50$, could be stalemated this way. For this reason, Mueller argues that decision-making costs should increase by a constant for $k \leq 50$, as depicted in Figure 2. Such a jump discontinuity, or "kink" in the decision cost function, will lead to a minimization of total costs at majority rule, which might explain why majority rule is so ubiquitous.

Even though Mueller makes an important step in the study of k-majority rules, there is no reason to assume that all voting thresholds less than $N/2$ must allow for contradictory proposals. For example, the U.S. Supreme Court requires only four of its nine justices to issue a writ of certiorari. Once the writ is issued by four justices, the writ cannot be recanted by four or five of the remaining justices, because the k-majority rule only applies to the status quo of no writ. Put differently, there are cases where $k \leq N/2$, avoid self-contradictions and Mueller's jump discontinuity is not applicable.

Other scholars have expanded on Buchanan and Tullock's original idea. Spindler (1990) argues that institutional framers should consider rent-seeking costs when evaluating the optimal k-majority rule. With rent-seeking costs included, the optimal k-majority rule might be near $k = 1$ or $k = N$. Guttman (1998) argues that k-majority rules should be judged on the Kaldor-Hicks criterion. Brennan and Hamlin (2000) analyze the optimal proportion of representatives in a democracy and replace external costs with agency loss. Their findings are similar to those of Buchanan and Tullock with the optimal k-majority rule somewhere in the middle. Finally, Mueller (2003, 103) makes an additional argument that vote cycling makes decision costs quadratic for k-majority rules greater than $N/2$. Such an adjustment implies that total costs are minimized around 64% of the eligible voters.[4] The common element in all of these works is that mathematics is used descriptively.

Dougherty and Edward (2011, 2004) attempt to rectify that shortcoming by formalizing external costs and decision costs and deducing the optimal k-majority rule. Using probabilistic models, they find that both the external cost and the decision cost functions are largely affected by the probability of passing proposals as determined by the binomial formu-

---

[4] His claim is based on work by Caplin and Nalebuff (1988) who show that voting rules with $k \geq .64(N)$ will not produce vote cycles.

la. Because the binomial formula produces a logit-type curve across various k-majority rules, proposals will almost certainly pass for small values of k and almost certainly fail for large values of k, with a steep decrease in between (see Figure 3.A for the probability of passage with $p_1 = .5$). Because the probability of passing proposals is at the heart of expected loss and expected decision cost functions, both functions are logit-type shapes as well. Figure 3.B shows the external cost and decision costs they derive for one set of parameters. External costs and decision costs are logit shaped in this figure because the probability of passage dictates the types of proposals that will pass and the number of proposals required for passage. Their results depend on 1) the probability of favoring and opposing the proposal, 2) whether the vote is for a single proposal or multiple proposals in a series, and 3) the rate at which more favorable proposals are created from one round to another, among other parameters. They find that if decision costs are negligible, the optimal k-majority rule can be a range of k-majorities near unanimity rule, as depicted by the external cost function in Figure 3.B. Because external costs are roughly zero for $70 \leq k \leq 100$, the optimal k-majority rule would be within this range. If decision costs are included, the optimal k-majority rule could be near majority rule.

3 One Proposal

For simplicity of exposition we will refer to a generic individual as "he." To focus our analysis, assume N individuals are trying to choose a k-majority rule for their town hall meeting. They want to know the expected external costs (and expected decision costs) from future votes under each k-majority rule. In this section, the future vote will be an up or down vote on a single proposal. Think of the up or down choice as accepting or rejecting a proposed constitution or some type of referendum that affects the N individuals only. All members will vote for either the status quo, q, or the proposal, x, in that future vote. No one votes "abstain," making the distinction between simple and absolute k-majority rules irrelevant (Dougherty and Edward 2004, 2011). Hence, a k-majority rule can be defined as follows.

Definition *k-majority rule*: proposal x defeats the status quo, q, by k-majority rule if and only if $\#yeas \geq k$, where $1 \leq k \leq N$; otherwise q is chosen.

In the case of a single proposal, each individual i has utility (gain or loss) from passage of the proposal and is a member of one of G groups. Different groups reflect different tendencies in preferences for "low tax," "high spending," etc. Members of group g have their utility drawn from a normal distribution with mean μ_g and standard deviation σ_g. Assuming $i \in g$ we denote this by

$$u_i \sim N(\mu_g, \sigma_g).$$

Randomly drawing utility reflects an individual's uncertainty about the proposal they will consider. For example, individuals may know that the future vote will be a tax proposal, but at the time that the k-majority rule has to be chosen they may not know the form of the proposal. Hence, they do not know whether they will gain or lose under the proposal. We also assume in this section that for all g, $\mu_g \in [-1, 1]$.

In addition, let e_i be a random error term, reflecting the imperfect information individual i has about his decision. For example, after a tax rule is proposed and individuals have time to consider it, individual i may still be uncertain about whether he would be better off or worse off under the tax rule and misjudge his actual utility gains from passage of the proposal. We allow such error to vary by group but require that for $i \in g$, $e_i \sim N(m_g, s_g)$. The randomness of e_i represents different individuals having different information and ability to evaluate tax rules, the mean tendency of which varies by group. Individual i *perceives* his utility from passage to be

$$u_i + e_i.$$

If i perceives his utility to be positive, he will favor the proposal. If he perceives it to be negative, he will favor the status quo. More precisely, voting is conducted with the function

$p_i = f(u_i + e_i)$:

$$p_i = \begin{cases} 0 & \text{if } u_i + e_i \leq 0 \\ 1 & u_i + e_i > 0 \end{cases}$$

where p_i is the probability individual i votes yes. Note, if i perceives a gain in a proposal (as indicated by the sum of his utility and error), he will certainly vote yes; if i perceives a loss, he will certainly vote no.

Example. Suppose an individual i comes from a group that tends to gain from a tax rule but members typically misinterpret the tax and expect to lose from it. We might have $u_i \sim N(.3, .2)$ and $e_i \sim N(-.4, .2)$. With these distributions for u_i and e_i, the expected value of $u_i + e_i$ will be $-.1$, and such a voter is likely to vote "no." However, since u_i and e_i are random, it could be that the generated values are $u_i = .2$ and $e_i = -.1$, in which case $u_i + e_i = .1$ and $p_i = 1$. Such a voter would thus vote "yes."

The community votes using yeas and nays as determined by p_i. Using the total yeas, they determine if the proposal passes under each k-majority rule, and then evaluate external costs based on the negative utility they expect. Note, voting and passage of the proposal are based on perceived utility $u_i + e_i$, whereas external costs are based only on values of u_i.[5]

We report three different measures of external costs (i.e. expected loss). First, "typical external costs" is the average of the losses incurred by all voters. The "worst group external costs" (resp. "best group external costs") is the average of the losses among all voters in the group with smallest (resp. largest) mean utility generated.

Example. Suppose $N = 3$ and we divided the voters into *two* groups. Group A: $u_1 = .8$, Group B: $u_2 = .2$. and $u_3 = -.1$. Thus, the loss for voters 1 and 2 would be 0, and the loss

---

[5] External cost is only incurred by voters who have negative utility when a proposal passes. When the proposal passes, external cost for a voter is $|u_i|$ if $u_i \leq 0$, and 0 if $u_i > 0$.

for voter 3 would be .1. Furthermore,

$$\text{typical external costs} = |0 + 0 + (-1)|/3 = .33;$$
$$\text{worst group external costs} = |0 + (-.1)|/2 = .05;$$
$$\text{best group external costs} = |0|/1 = 0.$$

The results for the worst group would be of particular interest for those who expect to be on the dissenting side of an issue, for example, those who believe a tax plan will be proposed that spends the proceeds on others but taxes them. We compute the expected losses from this framework using simulations, described in appendix A.

3.1 Results

The external costs for $N = 100$ individuals, $G = 3$ groups, and various group sizes are depicted in Figure 4. They reflect the losses an individual expects from various k-majority rules for each of the groups they might be in. In the first three frames, the utility of the three groups are drawn from $N(-.3, .2)$, $N(0, .2)$, and $N(.3, .2)$, respectively. The fourth frame is the same as the first, except the mean of the first and third groups are smaller, as they would be if the first and third group were more likely to oppose the proposal. For each group, we assume the error in judgement, is $e_i \sim N(0, .3)$.

In frame A, the three groups contain 35, 30, and 35 members respectively. The group means for drawing utility values, are symmetric around a mean utility of 0. This reflects a case where the same number of members will gain from the proposal, as will lose from the proposal. Thus, expected gains and losses should offset each other. If the individual had no prior knowledge of the group he was in *ex ante*, he should consider the mean loss depicted by the solid black line. The voter would expect external costs slightly greater than .1 for all $k < 45$, a sharp decline in external costs for $45 \leq k \leq 55$, and relatively zero external costs for $k > 55$. However, if the individual thought he would be a member of the first group, with expected utility of $-.3$, then his expected external costs would be greater, as depicted by the dotted line for the worst group. In contrast, if the individual thought he would be a member of the third group, with an expected utility of .3, then his expected external costs would be roughly zero for all k-majority rules because he would be very likely to gain from the passage of the proposal and not incur a loss under any k-majority rule. External costs for the best group are depicted by the horizontal, dot-dash line in Figure 4.A. As in Dougherty and Edward's previous work, there are are range of k-majority rules that minimize external costs ($k > 60$ in this case).

Similar stories are shown in Figure 4, frames B and C. In frame B, the three groups have 20, 30, and 50 members respectively, reflecting *more* individuals favoring the proposal than opposing it. With more individuals favoring the proposal, the height of the external cost function decreases for the typical individual and the steep decline in the external cost function begins further to the right. In frame C, there are 50, 30, 20 members of the three groups respectively, indicating that *less* individuals are expected to favor the proposal than oppose it. The mean loss is greater in this case than it is in frame A, but the number of yea votes in the system also decreases, pushing the location of the steep decline slightly to the left. If more individuals expected to have negative utility, the location of the steep decline could be even further to the left.

Finally, frame D has the same parameters as frame A, except the worst group has a expected utility of $-.5$ and the best group as an expected utility of .1. In terms of utility, this frame contains the same number of expected gainers and expected losers as frame A *ex ante*, but a third of the members of the best group will have a loss *ex post*. This explains why the height of the expected cost curve for the typical voter, the worst group, and the best group are all larger than those depicted in frame A. Because some members of the best group will have losses in utility and even more individuals will vote against a proposal they favor, there will be more nay votes in the system, which explains why the steep decline occurs for smaller k in frame D than it does in frame A.

Despite the paper's explicit use of different groups of voters, these results are fairly reminiscent of Dougherty and Edward's early work, with the external costs from each k-majority rule positive when the probability of passage is roughly one and zero when the probability of passage is roughly zero.

Perhaps the most interesting insight from all these cases, and several cases not depicted, is that regardless of which group an individual expects to be in (worst, best, or undetermined), external costs will decline to zero in roughly the same place. This is because the probability of passage dictates the plateaus of the external cost function. Over a certain range, the proposal will almost certainly pass and external costs will be positive. Over another range the proposal will almost certainly fail and external costs will be roughly zero – regardless of the number of groups or who is within them. Which group an individual expects to be in affects the height of their external cost curve, but, without decision costs, it should not affect the k-majority rule an individual finds optimal.

3.2 Decision Costs

At this point it might seem appropriate to include a discussion of decision costs. According to Buchanan and Tullock (1962), decision costs are the "time and effort . . . required to secure agreement" (p. 68). In the next section, we assume there is a cost for creating and considering proposals each round and that some k-majority rules will incur more decision costs because they require more rounds for an acceptable proposal to be created. The same argument does not apply to a single proposal. A single proposal prohibits any restructuring of the measure to make it more palatable. Hence, the time and effort needed to consider a single proposal, like a referendum, should be constant across all k-majority rules. This explains why decision costs are not included in Figure 4. With constant decision costs, the optimal k-majority rule is the one that minimizes external costs.

4 Multiple Proposals in a Series

We now extend our analysis to cases with multiple proposals voted on in a "successive" procedure. In a successive procedure, the initial status quo q_1 is paired against a proposal x_1 in round 1. If x_1 passes, voting ends. If x_1 fails, q_1 is paired against x_2 in round 2, and so on for a total of R rounds (see Figure 5). In a standard successive procedure, voting continues until the proposal passes or a specific round R is reached and the proposal either passes or fails.[6] In our study, voting continues until the proposal eventually passes, which we assure by improving the popularity of proposals in successive rounds. We also examine a few cases where some individuals are made better off and others are made worse off as the series proceeds. In these cases, we artificially stop voting as done in the standard successive procedure because no proposal will pass for larger k-majority rules. Successive procedures are used by legislatures in France, Germany, Greece, Spain, and other European countries, to name a few (Rasch 2000). It is also the multi-round procedure Tullock (1998, 86) considered for assemblies.

Each round of the series involves three steps. In the first step, a proposal is made extraneously. Voter i's perceived utility from passage of the proposal is then determined by draws from two distributions for each $i \in g$: the utility distribution $u_{i,1} \sim N(\mu_{g,1}, \sigma_g)$ and the error distribution $e_{i,1} \sim N(m_{g,1}, s_g)$, where $\mu_{g,1}$ and $m_{g,1}$ represent the mean utility and mean error for group g in round 1.[7] In the second step, the motion is discussed and decision costs are incurred. In the third step, the motion is put up for a vote. If the motion passes, the process ends. If the motion fails, then a new proposal is made with new values for $\mu_{g,r+1}$.

Continually improving proposals can be modeled in several ways. We assume that the expected utility of a group increases by an increment $\alpha_{g,r}$ each round; thus

$$\mu_{g,r+1} = \mu_{g,r} + \alpha_{g,r}, \text{ and } \alpha_{g,r} > 0.$$

---

[6] The final round, R may result from a limitation in the number of proposals or a time limit, etc.

[7] For the first round, we assume $-1 \leq \mu_{g,1} \leq 1$, but we relax this restriction for later rounds.

We consider two types of $\alpha_{g,r}$. The first is a constant, independent of r, reflecting cases where the expected utility of a group increases by the same amount each round. The second is a constant divided by r, ie. $\alpha_{g,r} = \frac{\alpha_{g,r}}{r}$, so that $\alpha_{g,r}$ increases at a decreasing rate. This models cases where it becomes more difficult to improve proposals as the rounds progress. We also consider cases where $\alpha_{g,r}$ is positive for some groups and negative for others, which require an artificial stopping point (i.e., a preset maximum number of rounds) because some larger k-majority rules will not pass a proposal in such cases.

To keep the analysis of decision costs simple, we assume each round imposes the same decision costs on all members of the assembly, $c > 0$. Decision costs are incurred in the second step. Voting rules which require many rounds to pass a proposal will generate more decision costs than voting rules which require few rounds.

Three assumptions should be made explicit:

(i) in each round, $\mu_{g,r}$ and $m_{g,r}$ are the same for all members of the same group, but they may vary between groups;

(ii) in each round, a voter's utility and error are independent of the utility and error of other voters in his group; and

(iii) across rounds a voter's utility is dependent upon his utility in the previous round.

Even though we allow for heterogeneous preferences and some probabilistic dependencies, our model does not relax the assumption of probabilistic independence completely.

4.1 Results: Typical Cases

We begin discussing our results using Figure 6. The figure presents external costs (thin solid lines), decision costs (thick solid lines), and total costs (dashed lines) for two sets of parameters. Both sets contain the same three groups. Group 1 has 35 members with an expected utility of $N(-.3, .2)$ in the first round; group 2 has 30 members with an expected utility of $N(0, .2)$ in the first round, and group 3 has 35 members with an expected utility of $N(.3, .2)$ in the first round. Each group also has an expected error $e_{i,r} \sim N(0, .3)$, which remains fixed across rounds. This reflects a fairly symmetric case with two groups offsetting gains and losses in the first round. The only difference between column A and column B is the change in expected utility across rounds. In column A, the mean expected utility of the three groups increases quickly by $\alpha = .1$ each round. In our simulations, all k-majority rules pass a proposal within eight rounds. In column B, the mean expected utility of the three groups increases slowly by $\alpha = .1/r$ each round. In our simulations, all k-majority rules pass a proposal within 60 rounds.

Each column reports the expected costs for the typical voter on top and the expected costs for the worst group below. The expected costs for the mean of the best group are not depicted. Decision costs are identical in the upper and lower frame of each column. Because the magnitude of external costs vary slightly, depending upon the perspective of a typical member or the perspective of a member of the worst group, total costs will vary slightly as well.

It may not be clear to the reader why external costs vary across k if we guarantee that a proposal will pass for all k. The answer is that smaller k-majority rules are more likely to pass proposals in earlier rounds, where $\mu_{g,r}$ are small and the utility generated is more likely to be negative. Larger k-majority rules are more likely to pass proposals in later rounds, where $\mu_{g,r}$ are large and the utility generated is less negative. The difference is not whether a k-majority rule allows a proposal to pass by the end of the series. In the cases depicted in Figure 6, it certainly will. The difference is whether a k-majority rule will allow a proposal to pass which appeals to a few voters or cause the assembly to wait until later rounds where proposals appeal to a larger proportion of voters.

Consider column A of Figure 6. In this case, the first round proposal typically passes for $k \leq 60$. For $k > 60$, expected utility increases quite rapidly between rounds, making all k-majority rules almost certain to pass a proposal within roughly eight rounds. Because the proposal typically passes in the first round for $k \leq 60$, the mean loss in Figure 6.A is the same as in Figure 4.A for $k \leq 60$. For $k > 60$, external costs decline much more slowly in Figure 6.A than in Figure 4.A, because $k > 60$ will pass a proposal in the multi-round case and each $k + 1$ will pass a proposal in a later round than k, producing a smaller loss.

In terms of decision costs, for $k \leq 60$ decision making decision costs equal the per round decision costs of $c = .01$ because the proposal passes in one round for these k. For $k > 60$, a proposal is quickly found that all individuals favor, explaining the slow increase in decision costs for $k > 60$. Because external costs exceed decision costs for most k, the sum of the two functions is minimized at $k = 95$ for the typical voter.

The outcome for the worst group is similar, except the external costs for a member of the worst group are much greater than the external costs for the typical voter. This explains the higher plateau in the bottom-left frame of Figure 6 compared to the top-left frame. Because the height is greater, the sum of the external cost and decision cost functions are minimized at the slightly larger value of $k = 99$. Whether an individual considers the choice from the perspective of the worst off group or just a member of the community affects the optimal k-majority rule, but surprisingly, not by much.

Column B of Figure 6 differs from column A in the rate at which the probability of favoring the proposal increases between rounds. In this case, external costs look almost identical to those in column A. This is largely because group sizes, initial utility, and expected error are the same in both columns. For $k \leq 60$, the proposal typically passes in the first round and decision costs are roughly .01 in both columns. The major difference between the two columns is in the decision costs for $k > 60$. Decision costs increase at an increasing rate for these k-majority rules because expected utility increases at a decreasing rate for $\alpha = .1/r$. Meaning, it takes increasingly more rounds to pass a proposal as k increases with this rate of growth. The inability to find proposals that favor an ever increasing threshold of voters, increases decision costs for larger k-majority rules and effectively decrease the optimal k-majority rule. In this case, total costs are minimized at $k = 82$ for the typical voter and at

$k = 88$ for the worst group, supporting Wicksell's ([1896] 1967) assertion that a k-majority "near" unanimity rule would be ideal.

A member of the best group, not depicted in Figure 6, should expect gains from any of the proposals made in any round. As a result, he would not expect to suffer external costs for any k-majority rule and the external cost function would be roughly zero for all k. Total costs would be minimized for this group, where decision costs are minimized. Namely, $k \leq 60$ in both columns.

Of course, the optimal k-majority rule depends on the per round decision costs, c. Figure 7 displays results for the same cases as shown in Figure 6, except the per round decision costs $c = .1$. As one might expect, decision costs are steeper and total costs are minimized at much smaller k-majority rules.[8]

For those interested in more analytical results, appendix B shows a more stylized, non-stochastic version of the model. This model demonstrates that the optimal k-majority rule largely depends on the per round decision costs c and the incremental improvement in the proposals α. It also helps the reader see understand the mechanism behind our model.

4.2 Results: Special Cases

Although external costs typically decline more linearly for larger k-majority rules in a series than in a single up or down vote, the decline can be abrupt. Figure 8 depicts two special cases. Both contain three groups: group 1 has 50 members with $u_{i,1} \sim N(-.5, .2)$; group 2 has 30 members with $u_{i,1} \sim N(0, .2)$; and group 3 has 20 members with $u_{i,1} \sim N(.1, .2)$.

---

[8] Buchanan and Tullock (1962) discuss the effect of preference homogeneity on the optimal k-majority rule. Our framework allows for several formalizations of preference homogeneity, such as changes in the relative size of each group and the distribution of the initial utility. In our model, the homogeneity of the group has little effect on the optimal k-majority rule independent of its effect on the round where a proposal passes and the losses incurred in those rounds — themes we have examined in the text. For example, three groups sized 30, 35, and 30 with initial utility $N(0, .2)$, $N(0, .2)$, $N(0, .2)$, and $\alpha = .01$ might produce larger external costs and smaller decision costs for $k > 65$ than the same groups with initial utility $N(-.9, .2)$, $N(0, .2)$, $N(.9, .2)$. But the difference is largely due to the the first set passing a proposal more quickly than the second set.

Column A reflects a case where each new proposal increasingly favors groups 2 and 3 across rounds, but increasingly goes against the interests of group 1 as the rounds progress. Error is distributed as previously. Because the proposal will certainly not pass for larger k-majority rules, for these parameters we artificially stopped the series at $R = 10,000$. The large number of rounds assures that a proposal passes for $k \leq 50$, but because utility is worsening for 50 members, a proposal will not pass for $k > 50$. The initial proposal typically passes for $k < 40$, making external costs for these k-majorities the same as for an up or down vote. They are zero for $k > 50$ because a proposal never passes in those cases. The interesting part of the external cost function is the slow "increase" in between, for roughly $k \in [40, 50]$. The increase occurs because group 1 (the worst group) is becoming increasingly worse off with each additional round, which causes an increase in their expected loss. Because external costs only measure losses, the gains of the other groups are not offsetting. Each increment of k from $k = 40$ to $k = 50$ requires one more favorable vote, which typically means more rounds to attain enough voters from groups 2 and 3 and more loss from group 1. Column A demonstrates an important exception to the common belief that external costs must decline monotonically as the k-majority rule increases. Considering external costs alone, a member of the worst off group, or the typical voter, would want the community to approve a proposal in an earlier round rather than a later round.

Adding decision costs accentuates the point. In Figure 8 column A, decision costs start at roughly .01 for those k-majorities where the proposal will pass in the first round, then they abruptly increase near $k = 50$ for those k-majorities that never pass a proposal.[9] As a result, the sum of external and decision costs will be minimized over the range $k \in [1, 37]$. The best hope for an individual who expects their utility from each additional proposal to worsen is to press for a k-majority rule that passes a proposal quickly or, better yet, avoid the collective agreement altogether. Demanding unanimity rule would not be the best option for them.

---

[9] For $k > 50$, decision costs equal $c(R) = .01(10,000) = 100$.

Figure 8, Column B illustrates another interesting case where error in judgement actually allows Pareto preferred proposals to arise that do not pass. In this case, $\alpha_1 = .001$ and $\alpha_2 = \alpha_3 = .01$, which improves the utility for the first group at a much slower rate than for the second and third groups. However, the most notable parameter in column B is the error term. We assume $e_{i,1} = N(-.1, .2)$ for group 1 and $e_{i,1} = N(0, .2)$ for groups 2 and 3, reflecting a case where the first group systematically misjudges proposals as against his interests. For example, people from South Carolina largely oppose federal revenue sharing, even though their state is one of the largest net beneficiaries of federal revenue sharing. In this case, there will be many rounds where $u_{i,r} > 0$ for everyone, but members of the first group will experience $u_{i,r} + e_{i,r} < 0$ and vote against the proposal.

The top-right frame of Figure 8 depicts results for the worst group. External costs are more of a logit-type shape in this case. They start flat for $k < 35$ and are roughly zero for $k > 80$. They decline abruptly in between because proposals shift from passing for smaller k-majority rules to failing for larger k-majority rule more abruptly. Of course, decision costs increase abruptly for $k = 50$ for similar reasons. Because decision costs for larger k-majority rules are substantially larger than external costs for smaller k-majority rules, the optimal k-majority rule from the perspective of the typical voter is any $k < 33$. The expected costs for the typical voter are not depicted, but they are fairly similar.

The bottom-right frame of Figure 8 depicts the proportion of rounds where a Pareto preferred alternative is proposed, but failed because members of group 1 erred in judgement. Diagnostics suggest that all k-majority rules will pass a proposal within 500 rounds, with $k < 50$ passing them in less than 25 rounds. As we should expect, larger k-majority rules pass a proposal in at least as many rounds as any smaller k-majority rule. But the lower-right frame of Figure 8 suggests something stronger. Larger k-majority rules fail to pass an even greater proportion of Pareto preferred proposals among their larger number of rounds. This is particularly problematic for $k > 85$. If an assembly adopts one of these rules, they may forgo a number of opportunities that make everyone better off.

5 Conclusion

External costs and decision costs are inextricably linked to the probability of passing proposals. The loss one expects to incur under a particular k-majority rule largely depends on the probability the k-majority rule passes a proposal with losses to one's group. The time and effort in decision costs largely depend on the same probability of passing proposals. In the framework used here, the probability of passing a proposal is not linear across k.

For a single proposal, the probability of passage takes on a logit-type shape, with passage almost certain for smaller k-majority rules and failure almost certain for larger k-majority rules. External costs have a similar shape for a single proposal. They will be positive for small k-majority rules because small k-majority rules can pass proposals that impose costs on others, and they will be roughly zero for large k-majority rules because large k-majority rules will not pass the same proposals. A range of k-majority rules with zero external costs imply that unanimity rule may not uniquely minimize external costs. Other k-majority rules, in the neighborhood of unanimity rule, may be equally adept by this criterion. If decision costs are negligible, or constant in an up or down vote, then society might have just as much justification adopting a number of supermajority rules as they would for adopting unanimity rule. Referenda on constitutions may be such a case.

For a series of proposals, the logit-type shape of the probability of passage translates directly into decision costs. Typically, decision costs are minimal over a range of smaller k-majority rules, they increase sharply over a middle range, and are large and almost constant for larger k-majority rules. A range of k-majority rules near $k = 1$ might minimize decision costs.

However, for a series of proposals, the logit-type shape of the probability of passage may not translate directly into the shape of the external cost function. Figure 8, column B shows a case where external costs follow that shape, but most of the cases we examined (both reported and unreported) show more gradual declines (as in Figure 6). This occurs because voting over a series of proposals is not a take it or leave it proposition. Larger k-majority rules can cause a body to wait for later rounds when more favorable proposals are developed with smaller losses.

Other factors that affect the optimal k-majority rule include the relative size of each group, the distribution of their initial utility, the amount of error in their judgement, how much subsequent proposals increase expected utility, and the decision costs incurred each round. Everything else equal, sequences that improve everyone's utility quickly will imply few decision costs for all k-majority rules. In those cases, large k-majorities, such as $k > .8(N)$ *may* be optimal. In sequences where utility improves more slowly cumulative decision costs will be greater for larger k-majority rules and the optimal k will be smaller. Some of the basic principles behind our model are illustrated by a stylized case in appendix B.

We attain more surprising results if subsequent proposals in a series make some individuals better off and other individuals worse off. In these cases, external costs may not decline monotonically as k increases. Instead, there can be regions where external costs actually increase as larger k-majority rules perpetuate voting and allow losing groups to accumulate greater losses. In such cases, the best way to protect individuals from collective decisions will be to set the k-majority rule at a small value or to forgo the collective agreement altogether.

Furthermore, if individuals systematically err in their judgement, they may not vote for proposals that advance their interests. However, this does not affect all k-majority rules equally. Communities voting under larger k-majority rules (near unanimity) will be substantially more likely to forgo a Pareto preferred proposal than a smaller k-majority rule. Again, the optimal k-majority rule may be in the neighborhood of $k = 1$.

Combined, our results suggest that different k-majority rules may be optimal in different circumstances. One parameter which drives that point home is the per round decision costs c. This parameter serves three functions in our model: 1) it measures the time and effort required to formulate a proposal each round, 2) it puts decision costs and external costs on the same scale, and 3) it provides a relative weighting between the two terms in the total cost function. The latter is clearly important because different individuals may value external costs and decision costs differently. For a series of proposals, there will always be a sufficiently large value of c that makes $k = 1$ optimal (perhaps among other k-majority rules in the neighborhood of $k = 1$) and there will always be a sufficiently small value of c that makes $k = N$ optimal (perhaps among other k-majority rules in the neighborhood of $k = N$). With different individuals valuing c differently, the optimal k-majority rule may vary among individuals in the same community. In which case, it might be very difficult to find a single k-majority rule that advances everyone's interests or to pre-suppose a voting rule for constitutional decision making. Even unanimity rule can favor some individuals more than others.

We conclude our paper with another oddity worth pondering. In our framework, actions are based on utility (gains and losses) plus error. These actions affected the passage of proposals and decision costs. However, the other part of our welfare analysis, external costs, is based entirely on expected losses. Potential gains in utility are not included in our normative analysis. We did this because Buchanan and Tullock (1962) seemed to focus on losses in the name of methodological individualism. But methodological individualism does not mean that individuals don't experience both the gains and losses from collective activity. Nor does it mean they will only consider the losses from entering into a collective agreement. Individuals might want to engage in collective activities because of the potential gains and avoid them because of the losses. In which case, they might want to chose a k-majority rule that maximizes their own expected utility rather than chose one that minimizes their expected losses. Future research might take this issue seriously and examine the optimal k-majority rule using a more traditional, utility based approach.

6 Appendix A

This appendix provides pseudo-code for the R programs used to run the simulations detailed in the paper.[10] The basic building block of a simulation is an *iteration*.[11] Each *iteration* begins with the same initial set of parameters, that are used for J iterations.[12] Averaging the external costs, decision costs, and total costs found in each iteration, approximates the expected values of the respective costs. It is those averages that are the model output presented in the paper.

Each iteration proceeds in the following order:

1. The per round decision costs, c, is assigned (sometimes a constant).
2. The group level parameters are assigned, including the size of each group, the mean and standard deviation of the utility distribution for each group ($N(\mu_g, \sigma_g)$), the mean and standard deviation of the error distribution for each group ($N(m_g, s_g)$), and the change in each group's mean utility ($\alpha_{g,r}$).
3. A *series of proposals*, indexed by r, is run according to a "successive procedure." For each round r:
    (a) Voter utilities and errors are generated based on their group level parameters:
        - An utility value, $u_{i,r}$ for each $i \in g$, representing utility from passage of the proposal in round r.

---

[10] The code needed to replicate all of the findings in the paper can be found on an anonymized dropbox folder at: `http://bit.ly/1clKkir`. Each file in the repository replicates one simulation presented in the paper. For example the file **three_goups_two_alt.R** replicates the simulation presented in Figure 4, and **series_results_typical.R**, replicates the results presented in Figure 6. There is one additional file, **user_external_decision_cost.R**, that allows a user to run a simulation based on parameters of their own choosing. Directions on running such simulations are provided in the file itself. The files may be downloaded individually or all the files for the entire paper may be downloaded as one zip file, **all_the_code.zip**.

[11] The vocabulary used in the psuedo-code includes the following. A **Round of Voting** refers to a case where the deliberative body votes up or down on a proposal. A **Series of Proposals** refers to the same deliberative body going through multiple *rounds of voting*, considering proposals one at a time, using a successive procedure. An **Iteration** is one pass through a entire *series of proposals*.

[12] We used $1,000$ iterations in the paper.

- An error value, $e_{i,r}$ for each $i \in g$, representing imperfect information.

(b) Perceived utility, $u_{i,r} + e_{i,r}$, for each voter is calculated and stored.

(c) $p_{i,r}$ is calculated for each voter based on their perceived utility.

$$p_{i,r} = \begin{cases} 0 & \text{if } u_{i,r} + e_{i,r} \leq 0 \\ 1 & u_{i,r} + e_{i,r} > 0 \end{cases}$$

(d) The number of *yea* votes is found. $yeas_r = \sum_i p_{i,r}$

(e) Whether or not the proposal passes for each of the k-majority rules in round r, is calculated and stored:

$$passes_{k,r} = \begin{cases} 1 & \text{if } yeas_r \geq k \\ 0 & \text{if otherwise} \end{cases} \quad \text{for } k = \{1, ..., N\}$$

(f) The potential external cost, $ec_{i,r}$, is calculated for each voter:[13]

$$ec_{i,r} = \begin{cases} 0 & \text{if } u_{i,r} > 0 \\ |u_{i,r}| & \text{if } u_{i,r} \leq 0 \end{cases}$$

(g) If all k-majority rules have passed a proposal, or the maximum allowed number of rounds have been reached, then no more proposals are considered and the series of proposals ends. If not:

   i. $\alpha_{g,r}$ is added to $\mu_{g,r}$ for all sub-groups.

   ii. The program advances to the next round starting at step 3a.

4. When the series of proposals ends for all k-majority rules, the following values are calculated.

---

[13] These costs are "potential" external cost, because an external cost is only incurred if a proposal passes. Although we keep track of potential external costs each round, we only report external costs for a k-majority rule in the round that a proposal passes (or the program stops due to a pre-specified number of rounds).

- R_k, which is the round each k-majority passed a proposal.[14]
- The decision cost incurred under each majority rule.

$$dc_k = c(R_k)$$

- The external costs for each k-majority rule:[15]

$$\begin{aligned} ecTypical_{k,R} &= mean(ec_{i,R}) \\ ecWorst_{k,R} &= max_g(mean(ec_{i,R,g})) \\ ecBest_{k,R} &= min_g(mean(ec_{i,R,g})) \end{aligned}$$

- Total costs for each k-majority rule:

$$\begin{aligned} tcTypical_k &= ecTypical_{k,R} + dc_k \\ tcWorst_k &= ecWorst_{k,R} + dc_k \\ tcWorst_k &= ecBest_{k,R} + dc_k \end{aligned}$$

5. This ends one *iteration* of the model.

6.1 Appendix B

In this section, we show analytical results for a specialized case. We assume the population consists of three groups, each one very homogeneous in its preferences. Group X has x voters, each with $u_{i,1} = \beta$; Group Y has y voters, each with $u_{i,1} = \gamma$; Group Z has z voters, each with $u_{i,1} = \delta$. Thus we are removing all the variance within groups and setting

---

[14] This is the lowest r, where $passes_{k,r} = 1$, for a given k-majority rule. If no proposal passed in any round, then the final round is used.

[15] Technically, we keep track of the $typicalEC_{k,r}$, $wosrtEC_{k,r}$, and $bestEC_{k,r}$ each round and update their values based on whether a proposal passed that round. We keep this process going for all k-majority rules across all rounds, then reference the results from round R_k for each k separately. This made writing of the code easier.

$\sigma_{i,r} = 0$ for all i, r. We also assume that there is no voter error, i.e., $e_{i,r} = 0$. These assumptions might model a parliamentary system with three parties each which has very uniform preferences, or a three person committee where members have different numbers of votes.

We further assume proposals improve by a constant α for all groups, the per round decision costs are c, and rounds continue until a proposal passes. We assume that $\beta > \gamma > \delta$, and also assume

$$\beta > 0 > \delta.$$

In what follows, we deduce the optimal voting rule k as a function of c, from the point of view of the worst group total costs (WGTC). Recall total costs are the sum of external costs and decision costs, i.e. $WGTC = WGEC + DC$. We will use the notation $WGTC_-$ to denote the $WGTC$ for $k \leq x$, $WGTC_0$ to denote the $WGTC$ for $k \in [x+1, x+y]$, and $WGTC_+$ to denote the $WGTC$ for $k > x + y$.

Case 1: $\gamma > 0$.

In this case, the proposal will pass in one round for $k \leq x+y$. For such k, $WGEC = -\delta$, and decision costs are $DC = c$, so

$$WGTC_- = WGTC_0 = -\delta + c. \tag{1}$$

For $k > x+y$, passage will require $r^* \geq 2$ rounds of voting, with r^* solving $(r^*-1)\alpha+\delta > 0$ and $(r^* - 2)\alpha + \delta < 0$.[16] Thus

$$r^* = \lfloor \frac{-\delta}{\alpha} \rfloor + 2. \tag{2}$$

For such k, we have $WGEC = 0$ and $DC = r^*c$, so

$$WGTC_+ = r^*c. \tag{3}$$

---

[16] With r^* rounds of voting, the proposal has only been improved $(r^* - 1)$ times, thus the factor $(r^* - 1)$ in this sentence.

We now compute c such that $WGTC_{-} = WGTC_{+}$. Solving for c using 1, 3, and 2 we get

$$c = c_{*} := \frac{-\beta}{r^{*} - 1} = \frac{-\beta}{1 + \lfloor\frac{-\delta}{\alpha}\rfloor}.$$

Thus for $c < c_{*}$, the optimal k-majority rule for the worst group will be any k in the interval $[x + y + 1, x + y + z]$. For $c > c_{*}$, the optimal k-majority rule for the worst group will be any k in the interval $[0, x + y]$.

Case 2: $\gamma < 0$.

Suppose first that $k \leq x$. Then the proposal will pass in the first round, and, as above, we have for such k

$$WGTC_{-} = -\delta + c. \tag{4}$$

For $x + y < k$, the measure will pass in r^{*} rounds, with r^{*} as in 2, and arguing as in Case 1,

$$WGTC_{+} = r^{*}c. \tag{5}$$

For $x < k \leq x + y$, the measure will pass in $r_{*} \geq 2$ rounds, where $(r_{*} - 1)\alpha + \gamma > 0$ and $(r_{*} - 2)\alpha + \gamma < 0$. Thus

$$r_{*} = \lfloor\frac{-\gamma}{\alpha}\rfloor + 2. \tag{6}$$

For such k, the $WGEC$ will be $-\delta - (r_{*} - 1)\alpha$, and $DC = r_{*}c$. Hence,

$$WGTC_{0} = -\delta - (r_{*} - 1)\alpha + r_{*}c. \tag{7}$$

It is intuitively clear that for c sufficiently large, the optimal k should be any $k \leq x$. To find the value of c where the optimal voting rule becomes $x < k \leq x + y$, we solve for c when $WGTC_{0} = WGTC_{-}$. By 7, 4, and 6 we get $-\gamma - (r_{*} - 1)w + mc = -\gamma + c$, hence

$$c = c^{*} := \alpha.$$

Thus for $c \geq c^*$, the optimal voting rule with be any $k \leq x$, while for c slightly under c^* the optimal voting rules will be $k \in [x+1, x+y]$. To find a lower bound on the c associated to this interval in k, we consider two sub-cases.

Case 2i: $\lfloor \frac{-\gamma}{\alpha} \rfloor = \lfloor \frac{-\delta}{\alpha} \rfloor$. In this case, $r_* = r^*$, and the number of rounds for $k \in [x+y+1, x+y+z]$ will be the same as the number of rounds for $k \in [x+1, x+y]$, and $WGTC_0 = WGTC_+$ for all c. Thus in this case, for any $c < \alpha$, all voting rules $k \in [x+1, x+y+z]$ will minimize the total costs incurred by the worst group.

Case 2ii: $\lfloor \frac{-\gamma}{\alpha} \rfloor < \lfloor \frac{-\delta}{\alpha} \rfloor$. In this case, we solve for $WGTC_0 = WGTC_+$. By 7 and 5, we get $r^* c = -\delta - (r_* - 1)\alpha + r_* c$ hence

$$c = c_* := \frac{\delta - (r_* - 1)\alpha}{r^* - r_*} = \frac{-\delta - (\lfloor \frac{-\gamma}{\alpha} \rfloor + 1)\alpha}{\lfloor \frac{-\delta}{\alpha} \rfloor - \lfloor \frac{-\gamma}{\alpha} \rfloor}$$

Thus for $c \leq c_*$, the optimal k-majority rule for the worst group will be any k in the set $[x+y+1, x+y+z]$. For $c \in [c_*, \alpha]$, the the optimal k-majority rule for the worst group will be any k in the set $[x+1, x+y]$.

Case 3: $\gamma = 0$. This case is computed similarly, with the details left to the reader.

Remarks

1. The numbers c_*, c^* are critical in the sense that as cost c crosses these numbers, the optimal voting rule changes. For this population, we are able to fully determine c_*, c^* if we know the the initial gains from passage of the proposal, γ, δ, of members of Y, Z, together with the incremental improvement in the proposals.

2. It is interesting to note that the optimal k does not depend on the population sizes x, y, z, nor on the initial utility from passage of the proposal, β, for the group X.

3. Because $c^* = \alpha$ in Case 2, one simple conclusion follows. If Group Y (the middling group) is initially opposed to the measure, the optimal voting rule will be $k \leq x$ if and

only if per round decision costs exceeds the incremental improvement in the proposal (i.e. $c > \alpha$).

4. The methods above can be used to determine the optimal voting rule k, as a function of cost c, from the perspective of the *typical* voter. However, in this case, the critical cost value c^* depends on y and z.

7 References

Brennan, Geoffrey and Hamlin, Alan. 2000. *Democratic Devices and Desires.* New York: Cambridge University Press.

Caplin, Andrew and Nalebuff, Barry. 1988. “On 64%-Majority Rule,” *Econometrica* 56(4): 787-814.

Dougherty, Keith L. and Edward, Julian. 2012. “Voting for Pareto Optimality: a multidimensional analysis,” *Public Choice* 151(3): 655-78.

Dougherty, Keith L. and Edward, Julian. 2011. *The Calculus of Consent and Constitutional Design.* New York: Springer.

Dougherty, Keith L. and Edward, Julian. 2005. “A Nonequilibrium Analysis of Unanimity Rule, Majority Rule, and Pareto,” *Economic Inquiry* 43(4): 855-64.

Dougherty, Keith L. and Edward, Julian. 2004. “The Pareto Efficiency and Expected Costs of K-Majority Rules,” *Politics Philosophy and Economics* 3(2): 161-89.

Guttman, Joel M. 1998. “Unanimity and Majority Rule: The Calculus of Consent Reconsidered,” *European Journal of Political Economy* 14(2): 189-207.

Heckelman, Jac C. and Dougherty, Keith L. 2010. “Majority Rule versus Supermajority Rules: Their Effects on Narrow and Broad Taxes,” *Public Finance Review* 38(6): 738-61.

Mueller, Dennis. 2003. *Public Choice III.* New York: Cambridge University Press.

Rasch, Bjorn E. 2000. “Parliamentary Floor Voting Procedures and Agenda Setting in Europe,” *Legislative Studies Quarterly* 25(1): 3–23.

Spindler, Zane A. 1990. “Constitutional Design for a Rent-Seeking Society: Voting Rule Choice,” *Constitutional Political Economy* 1(3): 73–82.

Tullock, Gordon. 1998. *On Voting: A Public Choice Approach.* Northampton, MA: Edward Elgar Publishing.

Wicksell Knut [1896] 1967. "New Principle of Just Taxation," in Richard A. Musgrave and Alan T. Peacock, eds. *Classics in the Theory of Public Finance.* Macmillan, pp. 72-118.

Figure 1: Buchanan and Tullock's External Costs and Decision Costs

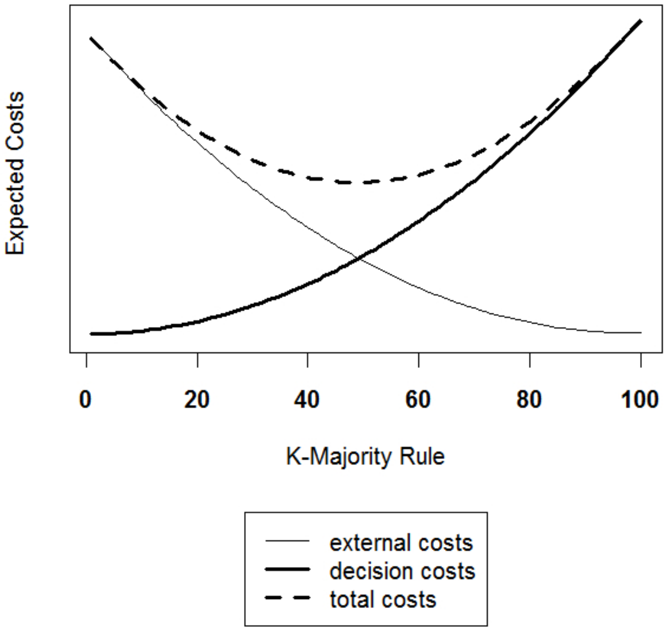

Figure 2: Mueller's Kink in the Decision Cost Function

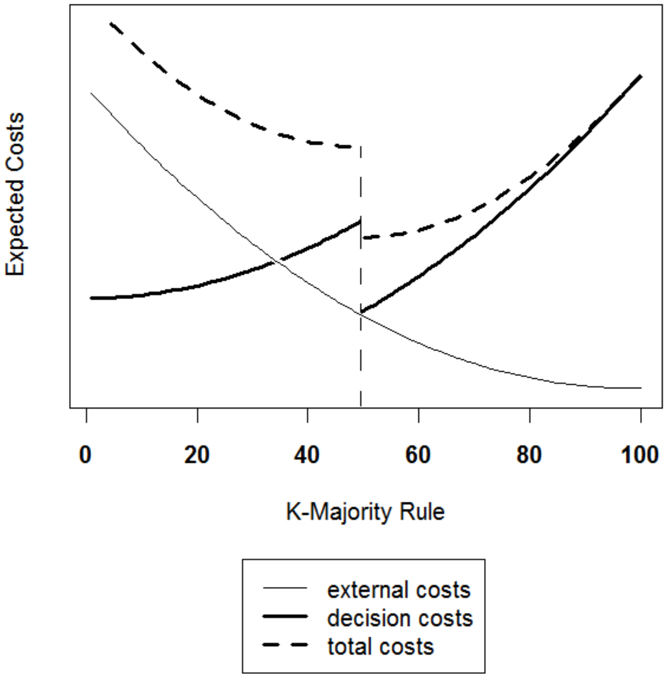

Figure 3: Dougherty and Edward's External Costs and Decision Costs

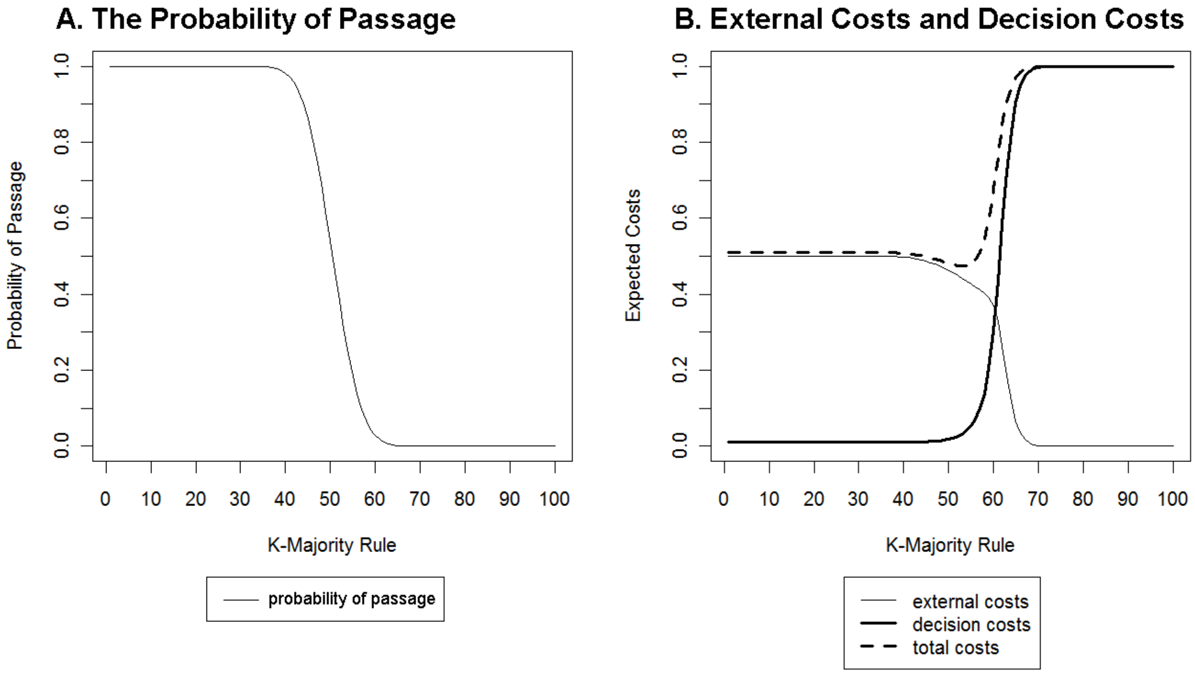

Notes: Frame A presents the probability of passing proposals assuming individuals are equally likely to favor the status quo and the proposal. Frame B presents external costs and decision costs assuming multiple proposals are voted on in a successive procedure, individuals are equally likely to favor a proposal and the status quo in the first round, per round decision costs of .01, and the probability of favoring the proposal increasing at the rate of $.001/r$ for each round r (see Dougherty and Edward 2001, 68).

Figure 4: External Costs from a Two Alternative Vote, Three Groups

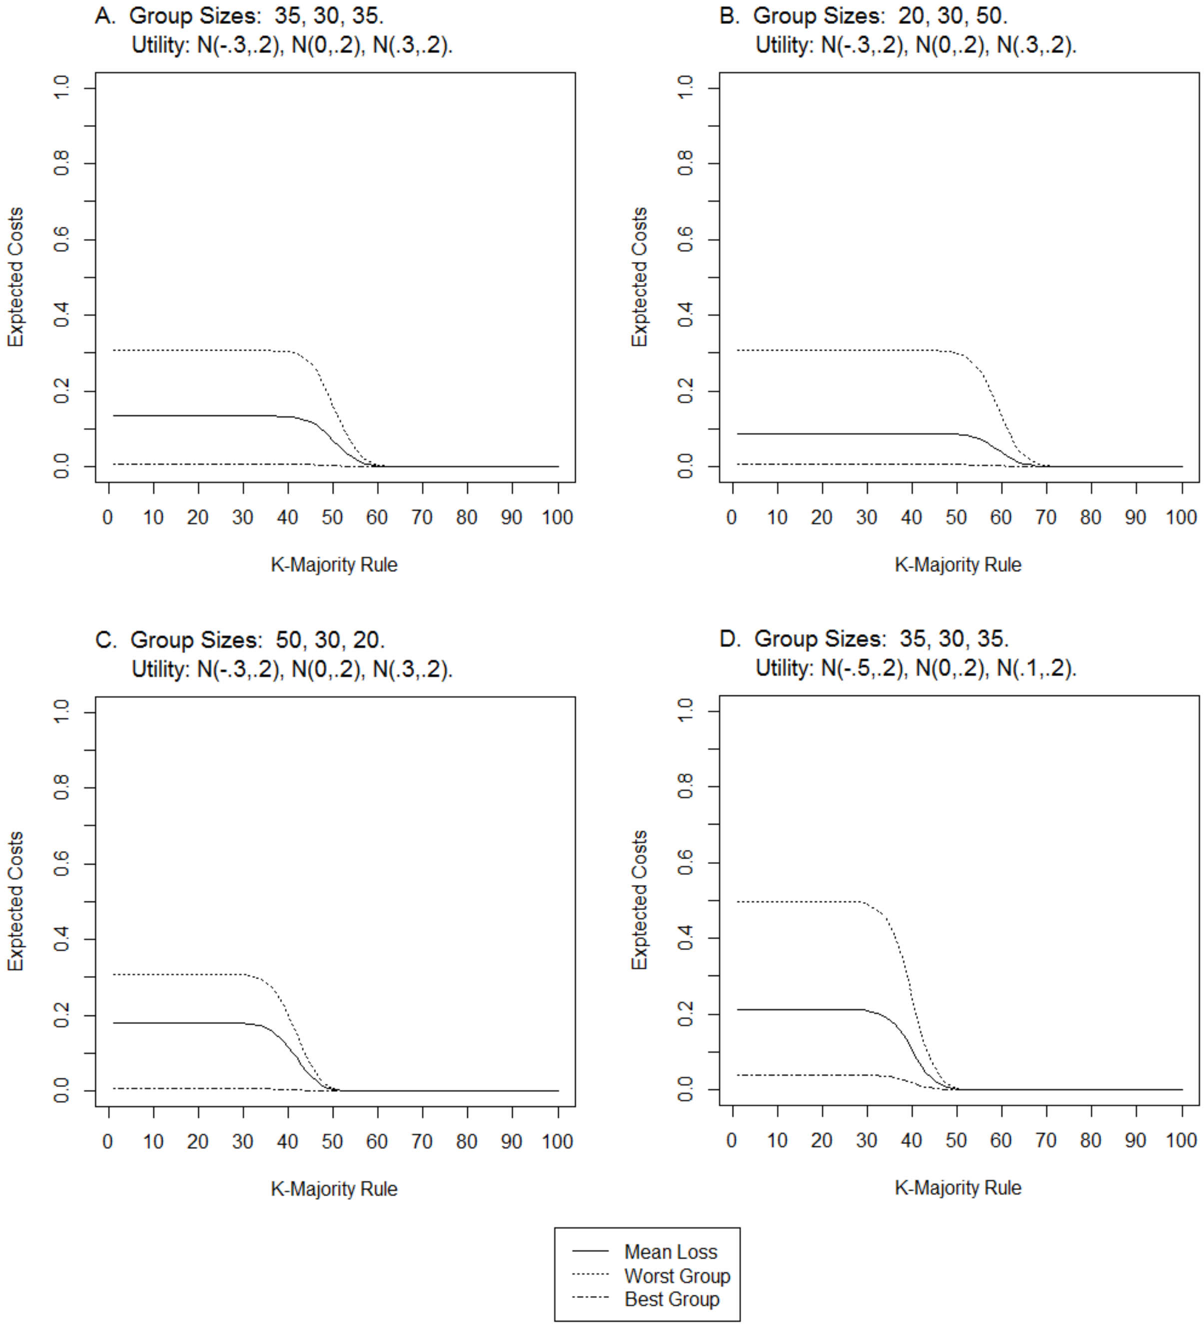

Note: For each case, $e_i \sim N(0, .3)$, $N = 100$.

Figure 5: A Successive Procedure

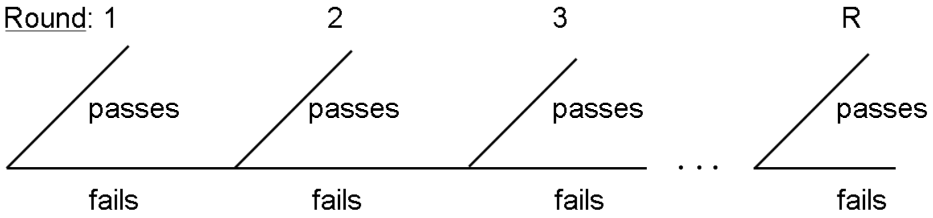

Figure 6: Expected Costs from a Series of Proposals: Typical Cases

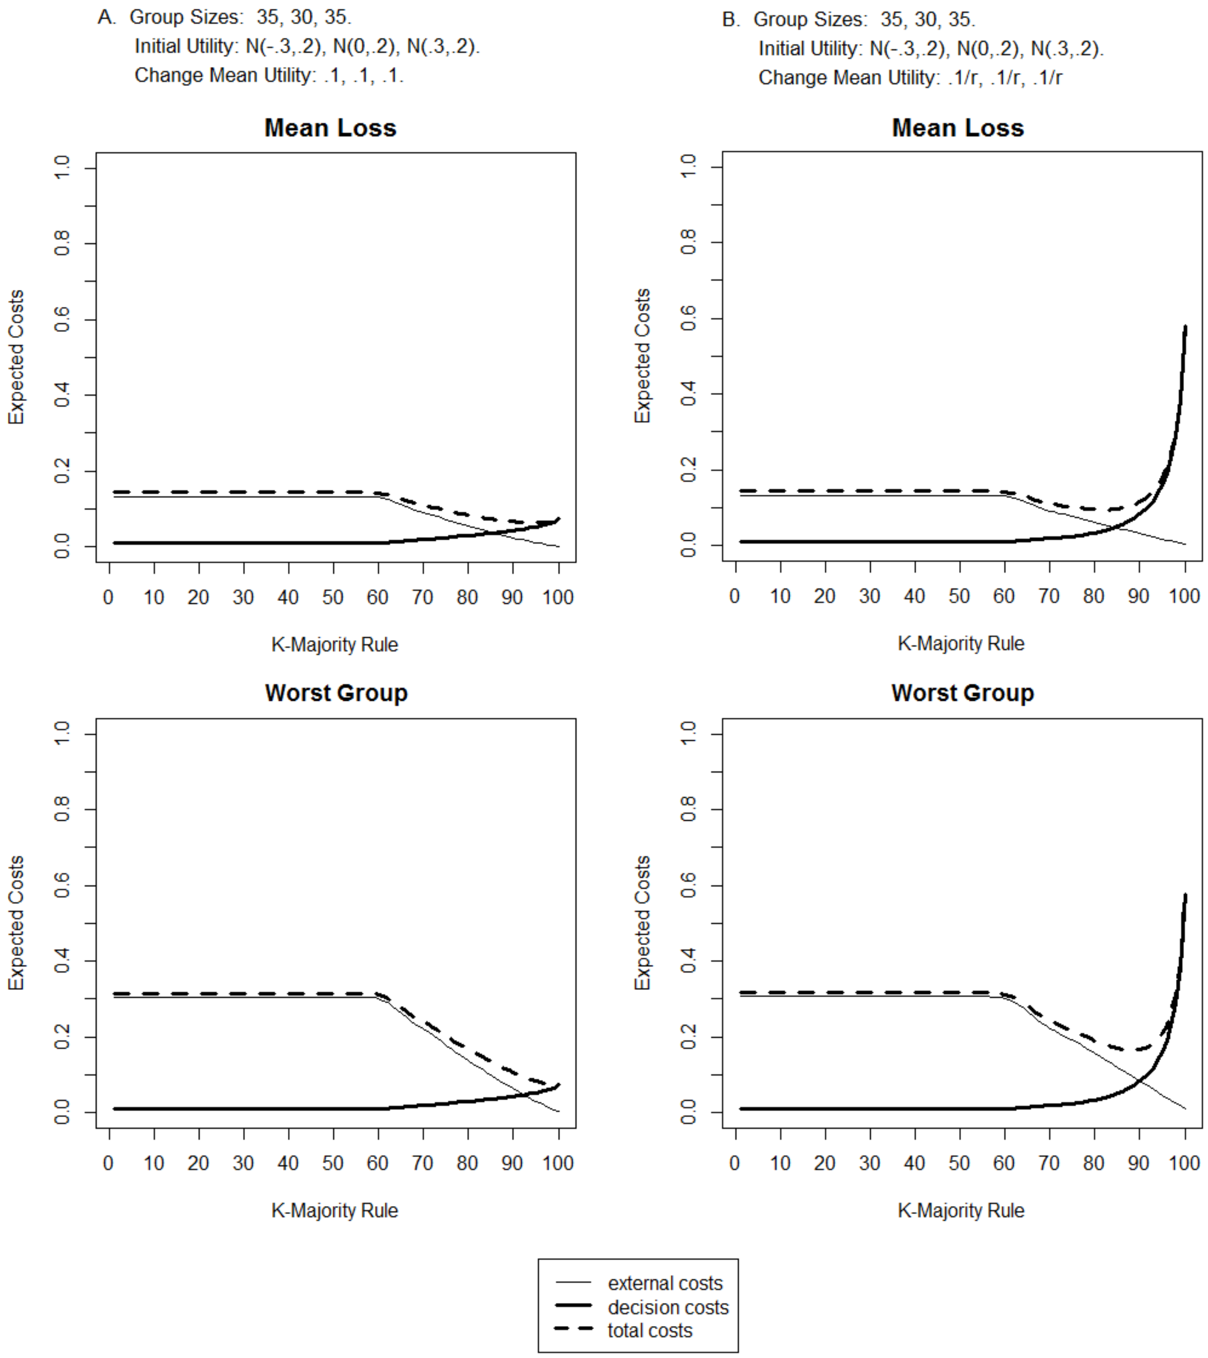

A. Group Sizes: 35, 30, 35.
Initial Utility: N(-.3,.2), N(0,.2), N(.3,.2).
Change Mean Utility: .1, .1, .1.

B. Group Sizes: 35, 30, 35.
Initial Utility: N(-.3,.2), N(0,.2), N(.3,.2).
Change Mean Utility: .1/r, .1/r, .1/r

Note: For each case, $c = .01$, $e_{i,r} \sim N(0, .3)$, and $N = 100$.

Figure 7: Expected Costs from a Series of Proposals: Larger c

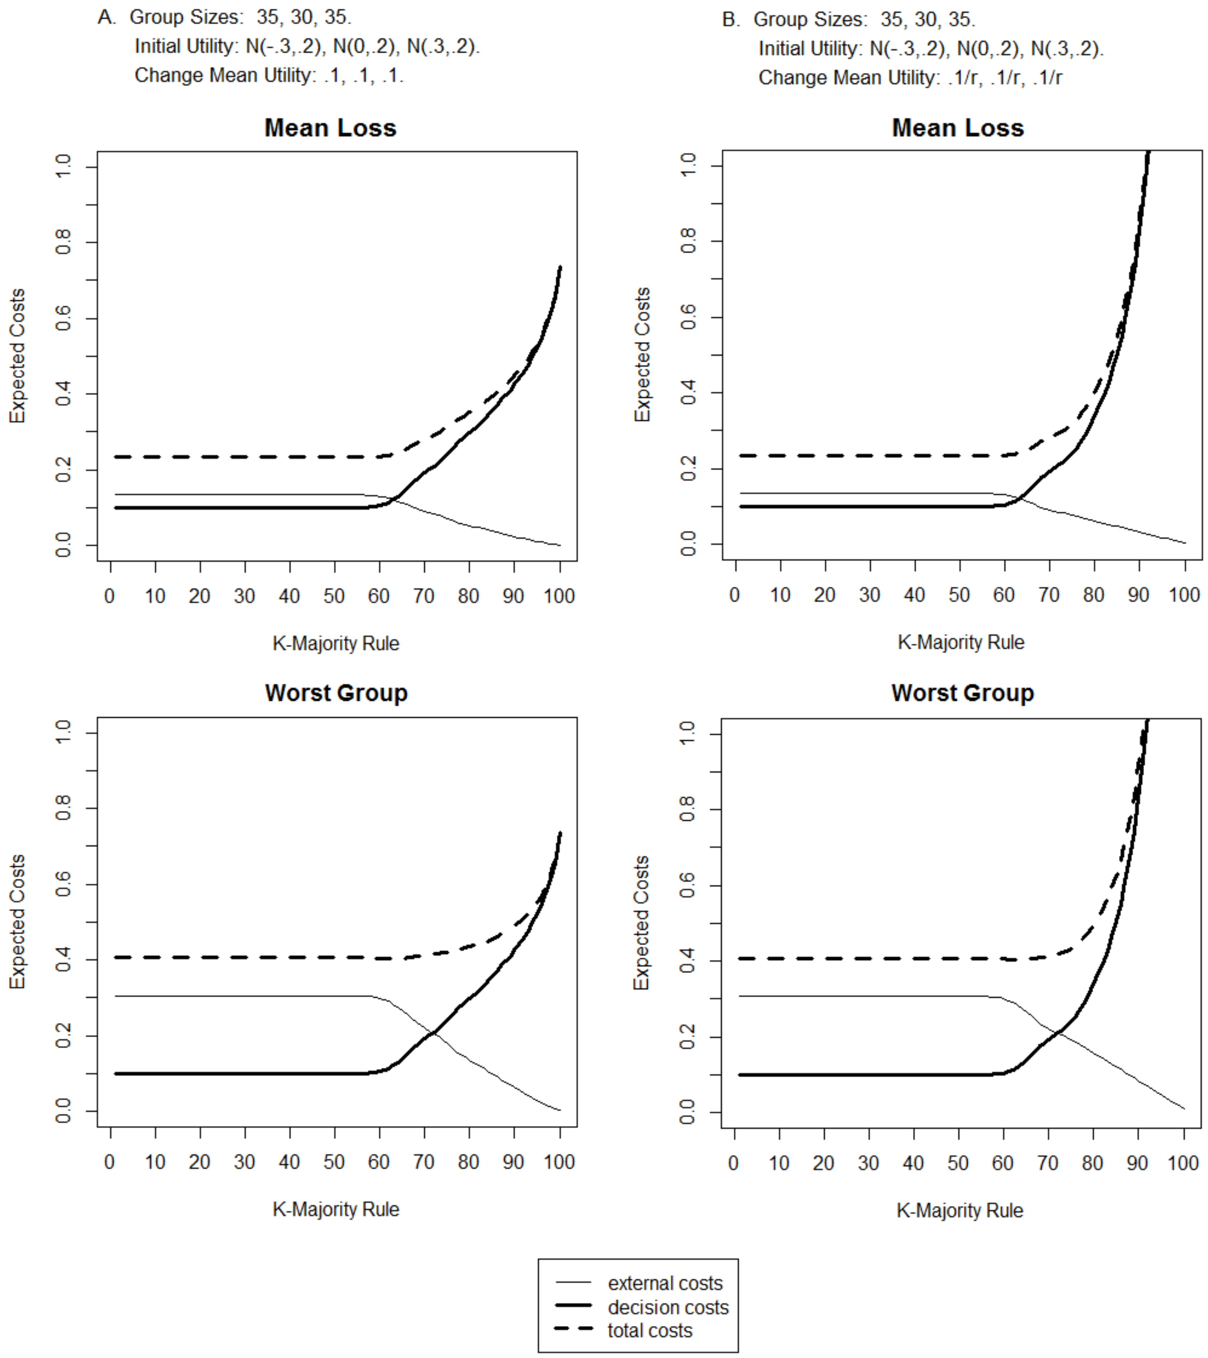

Note: For each case, $c = .1$, $e_{i,r} \sim N(0, .3)$, and $N = 100$.

Figure 8: Expected Costs from a Series of Proposals: Special Cases

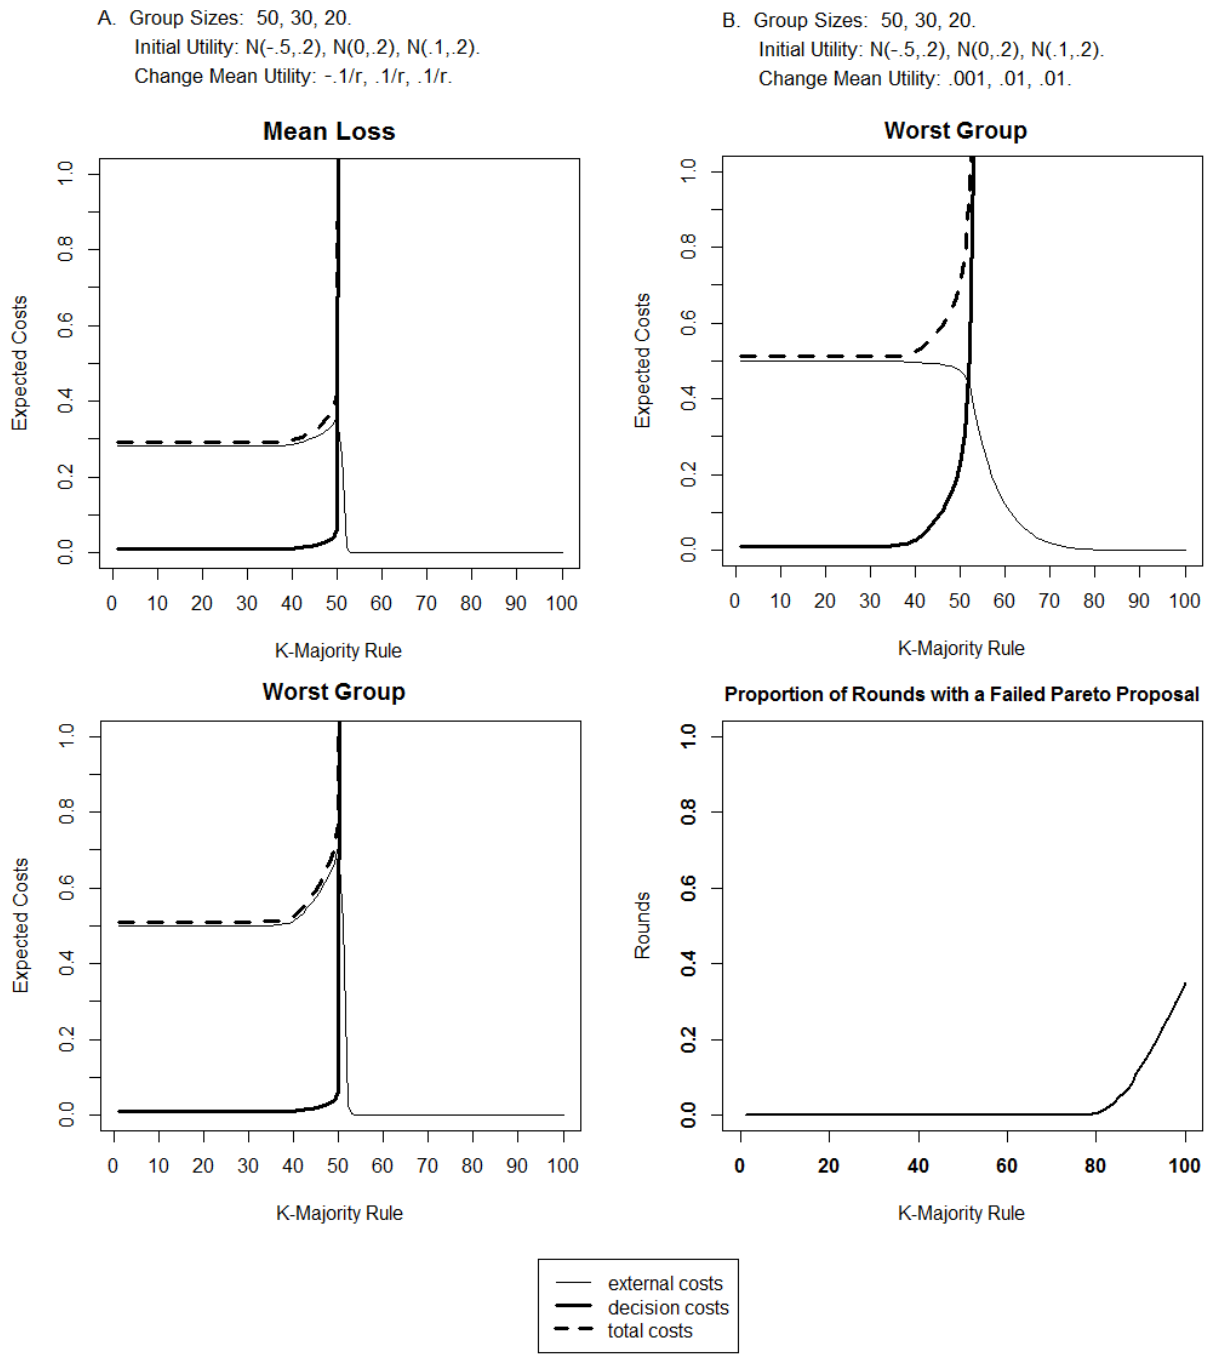

A. Group Sizes: 50, 30, 20.
Initial Utility: N(-.5,.2), N(0,.2), N(.1,.2).
Change Mean Utility: -.1/r, .1/r, .1/r.

B. Group Sizes: 50, 30, 20.
Initial Utility: N(-.5,.2), N(0,.2), N(.1,.2).
Change Mean Utility: .001, .01, .01.

Note: $c = .01$, $N = 100$. In frame A, for all $g, e_{i,r} \sim N(0, .3)$. In frame B, $e_{i,r} \sim N(-1, .2)$ for group 1 and $e_{i,r} \sim N(0, .2)$ for groups 2 and 3.