UNIVERSITY OF Southampton
School of Social Sciences

Economics Department
University of Southampton
Southampton SO17 1BJ, UK

Discussion Papers in
Economics and Econometrics

Information acquisition in citizen and representative assemblies

Maksymilian Kwiek (University of Southampton)

No. 2002

This paper is available on our website
http://www.southampton.ac.uk/socsci/economics/research/papers

ISSN 0966-4246

Information acquisition in citizen and representative assemblies

Maksymilian Kwiek*
University of Southampton

July 30, 2020

Abstract

A Representative Assembly (RA) consists of representatives elected by citizens. A Citizen Assembly (CA) consists of randomly selected citizens. I study the normative performance of these two institutions in a game-theoretic model with inattentive voters. The key question is which assembly creates incentives for voters to learn about the decision problem. Because I restrict how the two assemblies affect the probability of being pivotal, I can focus on issues other than the classical paradox of voting. The reason why the RA is inferior to the CA is that it forces the voters to learn about politics instead of policies. Since only policies matter, this is inefficient. The RA may be superior if the voter's preferences are highly correlated with those of the candidate, but a hybrid CA model (with candidates relegated to mere experts) is not worse and sometimes strictly better. I discuss the role of media in electoral campaigns and a number of other extensions.

JEL codes: D02, D71, D72, D83

Keywords: citizen assembly, sortition, voting, information acquisition, mutual information, entropy.

*Department of Economics, University of Southampton, Southampton, SO17 1BJ, United Kingdom. Email: M.Kwiek@soton.ac.uk.

1 Introduction

This paper investigates how two alternative forms of assembly affect voters' incentives to acquire necessary information. The first institution is a representative assembly (RA) which is composed of representatives elected by the voters in a general election. The other one is a citizen assembly (CA) in which it is a randomly selected subset of all citizens who makes a decision. The random selection process is sometimes referred to as sortition or selection by lot.

Representatives are used by the overwhelming majority of modern democracies. The United States Congress, Senate, and President are all elected, so are the House of Commons in the United Kingdom, the National Assembly and the President in France, and many other legislatures and executives around the world.[1] This model emerged in the late eighteenth and early nineteenth centuries (Manin 1997), although the House of Commons in the UK and the Sejm in the Polish-Lithuanian Commonwealth predate this period.

In response to the recent perceived crisis of democracy, many proposals have been put forward to improve the functioning of the democratic process (Smith 2005, 2009). The CA is one such innovation. This type of assembly has been proposed and used to solve intractable collective decision problems in a number of cases around the world. The most prominent examples include the 2004 Electoral Reform assembly in British Columbia, the 2016 Irish Constitution assembly, as well as many local government initiatives in the UK, Poland, and other places.[2]

The CA is not a modern invention. Athenian democracy used sortition to select officials and councils. In some Italian cities of early Renaissance, the ballot was used in the fabulously complicated process of selecting city officials. Florence and Venice are often-cited examples, but this form of government has spread as far as the Iberian peninsula (Manin 1997).

I present game-theoretic models of CA and RA and compare their equilibria. The ultimate goal is to address the normative question: which of the two assemblies is likely to reach a correct decision more frequently, and whether the answer to this question depends on some interpretable parameters. Since these two types of assembly differ along many dimensions, one of the preliminary challenges is to find a way to make this comparison meaningful, and this is where the methodological contribution of this study lies.

The main hypothesis is the following: the CA will perform better than the RA, as long as candidates' and voters' favorite policies are ex-ante poorly cor-

[1]A prominent counterexample is Switzerland which often uses referendums.

[2]See, for example, Paulis et al. (2020). There are also constantly updated databases such as https://participedia.net.

related with each other, that is, when their preferences might be misaligned. The key observation is that the voter's information acquisition problem in the CA involves only learning own preferences, while the RA requires learning own preferences *and* the preferences of the candidates, to see whether they match. Since these preferences are uncorrelated, this is a harder problem. Mnemonically, we could say that the RA forces the voter to learn about both politics (candidates' preferences) and policies (own preferences), while the CA requires only the latter. Since only policies matter for efficiency, the RA is inferior.[3,4]

The RA may have an advantage too. If candidates' and voters' favorite policies are ex-ante highly correlated then the voter can simply vote for the candidate, knowing that their preferences are aligned, even if she actually does not know what she or her candidate prefers. Therefore, the voter can bypass the need to acquire information. Since pure CA does not have such a mechanism, it may be inferior. This effect is likely to be strong if the cost of information acquisition is high.

These two sides of the main result echo the two views of democracy that clashed over the centuries (Manin 1997). According to the principle of distinction, the democratic rulers should be distinct from the ruled, for example, better educated, more eloquent, or wealthier. According to the principle of resemblance, they should be like the general public, share their circumstances and sentiments. The model below clarifies that we should adhere to the principle of distinction in the sense of better information—the rulers should have as best information as possible to be able to make correct decisions. However, we should adhere to the principle of resemblance as far as the preferences are concerned—if the rulers' preferences are different than those of the ruled then the agency problem would be too severe. In short, the decision-makers should have (i) the same preferences and (ii) better information than the general public. The RA achieves that if the preferences of informed candidates and the uninformed voters are aligned, otherwise, the CA is a better mechanism.

Apart from the pure CA and pure RA, I also consider a hybrid CA model.

[3]This study has rational and consequentialist voters who treat politics instrumentally, but one could contemplate a behavioral model in which voters derive intrinsic (dis)utility from following politics. For example, if politics has an entertainment value, then having RA could eliminate learning about policies, perhaps entirely.

[4]Current studies of collective information acquisition have voters learning either about the policy issues (for example, Ben-Yashar & Nitzan 2001, Mukhopadhaya 2003, Persico 2004, Martinelli 2007, Gerardi & Yariv 2008, Koriyama & Szentes 2009, Oliveros 2013) or about the candidates to be elected, that is politics (Matějka & Tabellini 2017). My model of RA requires both.

Members of this assembly may select one of the alternatives (as in the pure CA), but it also has an additional option to follow the recommendation of the expert advisers (the same informed agents who play the role of the candidates in the pure RA). Since these experts are informed, following their advice may be as advantageous as in the pure RA. Clearly, this hybrid CA model is not worse than any of the pure assemblies, but, for low levels of the information acquisition cost, it is *strictly* better than either of them. Namely, adding such experts to the CA improves it, even when the pure CA is better than the pure RA in which the same experts can be selected as representatives.[5]

More generally, the effect of information acquisition technology is nuanced, all depending on what defines a "poor" technology. Interpreted as a higher cost of information acquisition, poorer technology may favor RAs, as explained above. But I also argue that if poor technology means the lack of access to good punditry and journalism then it may make it harder for the voters to learn about candidates in the RA, which would favor the CA. One of the contributions of this paper is to emphasize the role of frictions in information transmission by the mass media.

Simplifying assumptions

Many important factors are assumed away in the analysis below. The rest of this introduction lists these simplifications and explains why they are made.

Democratic institutions differ along many dimensions: Table 1 captures two of them. The columns show the size of the decision-making group of citizens, and so they also represent the average probability of being pivotal. The rows show the distance between the citizen and the ultimate decision-maker, so they measure how direct democracy is.

Four cases emerge: the referendum is when the full population decides directly without any intermediaries. The general election of representatives is an example of indirect democracy, like presidential elections or parliamentary system. The citizen assembly is when the final decision is taken by a small subset of the population. The fourth combination depicts a hypothetical democratic institution, in which a citizen assembly selects representatives who will make the final decision—here called an "electoral college as a CA".

This study ignores the fact that the probability of being pivotal may be endogenous, that is, it conflates the columns of Table 1.

[5]The debate shifts to a discussion of the hybrid system which draws advantages from both pure mechanisms. For example, the collection of articles edited by Gastil & Wright (2019) contemplates a bicameral system with one chamber working as a RA and the other one as a CA.

	Full population	Random subset of population
Direct	Referendum	CA
Indirect	General Election (RA)	Electoral College as CA

Table 1: Possible mechanisms

Assumption 1. Probability of being pivotal is exogenously constant between the mechanisms.

This is not because this aspect is not important. To the contrary, perhaps the most powerful practical argument in favor of the CA over the RA is that the former allows for the decision to be made in a small committee and, therefore, it incentivizes information acquisition. By contrast, the RA involves a general election, where direct incentives to learn about the issues and candidates' positions are virtually nonexistent. In other words, the RA suffers more than the CA from a version of the Downsian paradox of voting. The reason I do not focus on this issue is that it is already covered in the existing literature.[6] This study focuses on the distinction illustrated by the rows of Table 1. In fact, much of the analysis below can be undertaken even if we assume that there is a single voter in the population.[7]

Another important issue that I ignore is how to select citizens into the CA. For example, should the CA membership be completely random, or whether it should be demographically representative? This appears to cause some concern among many practitioners as problems of this nature may lead to the system lacking legitimacy. For example, an unrepresentative sub-sample of the population may be selected into a CA by pure bad luck. I by-pass this issue by assuming that voters' preferences are common across the population and so the composition of the CA does not matter.[8]

Assumption 2. Voters have the same preferences.

Finally, there is a lot of practical effort that focuses on how to organize CAs so that their internal proceedings and deliberations are effective. Expert

[6]The argument can be traced to Downs (1957). More recently, Ben-Yashar & Nitzan (2001), Mukhopadhaya (2003), Persico (2004), Martinelli (2007), Gerardi & Yariv (2008), and Koriyama & Szentes (2009) investigate how the size of the committee matters in creating incentives for information acquisition. Relevant experimental studies include Bhattacharya et al. (2017) and references within.

[7]Thus, for much of this paper, the referendums and citizen assemblies are formally indistinguishable.

[8]However, Bardhi & Bobkova (2020) show that various sortition methods may produce different results even if the voters have common preferences but different local knowledge.

advice should be heard, dispersed information aggregated and creative solutions proposed, features that are sometimes difficult to achieve in adversarial debates among career politicians in RAs. However, I am not studying deliberation as such, not even the voting rules that the CA could use. Specifically, I assume that

Assumption 3. The Citizen Assembly consists of one member.

Sometimes, a CA is proposed to break a political deadlock and work out a general direction of change, whereby the set of alternatives is not prescribed. In contrast to this, my model focuses on a simple binary agenda, therefore there is no need to address the question of who sets the agenda.

The rest of the paper proceeds as follows. Section 2 starts with a relatively simple decision problem in the pure CA—the single citizen studies the decision problem and gives her final verdict. The reader should pay attention to how information acquisition is modeled by means of mutual information, as this technique will be used extensively in the more complicated pure RA and hybrid models. The model of pure RA is presented in section 3. The main battery of results is presented in section 4. Section 5 shows the hybrid model, and section 6 discusses further extensions.

In terms of the method of analysis and presentation, I focus mostly on illustrative but representative simulations.

2 Model of Citizen Assembly

Consider a society of $n \geq 1$ citizens who have to collectively select one of two alternatives: a status quo or a reform. The fundamental assumption of this study is that the citizens do not know which of the two outcomes they prefer. Technically, I assume that there is a binary random variable X which describes whether a citizen prefers the reform to occur. If $x = 0$ then the voter prefers the status quo, if $x = 1$ then the voter prefers reform. The realization of this random variable is unknown to her. In line with Assumption 2, I assume that the single X describes the preferences of each of n citizens. Let $g_x(x)$ be the probability of x; assume that this distribution is uniform, $g_x(0) = g_x(1) = 1/2$.

The decision is taken by the CA, which, by Assumption 3, consists of a single randomly selected citizen. That member decides in favor of the status quo or the reform, which are denoted, respectively, by $c = 0$ and $c = 1$.

The decision problem of the voter (CA member) can be understood as consisting of two stages: in the information acquisition phase, the voter selects a costly precision of some signal of x, and then in the decision phase, she selects c as a function of the realized signal. In the terminology of Matějka & McKay (2015), this is a state-signal-action model. However, I adopt what they call a state-action model, whereby the decision-maker simply selects a conditional probability distribution $h(c|x)$, where the choice of the signal is implicit. This defines a random variable C, the marginal distribution of which is $\tilde{h}(c) = \sum_{x'} h(c|x') g_x(x')$.

Citizen's ex-ante preferences are described by a disutility function that consists of two elements. Firstly, she suffers a loss of 1 if the alternative selected is not the one that is preferred, that is, the loss function is $L(x,c) = 1$ if $x \neq c$, and 0 otherwise. The second element is the cost of acquiring information. This will be captured by mutual information between random variables C and X.

$$I(X;C) = \sum_x g_x(x) \sum_c h(c|x) \log_2 \frac{h(c|x)}{\sum_{x'} h(c|x') g_x(x')}$$

Mutual information has the following interpretation: it is the number of bits required on average in order to communicate action c characterized by $h(c|x)$, when the source x is characterized by $g_x(x)$. For example, it can be easily verified that if the voter does not learn anything about x and, consequently, c is independent of x, then mutual information is zero. In the other polar case, if c represents x perfectly ($h(c|x) = 1$ if and only if $c = x$), then mutual information is 1, which indicates that this strategy requires precisely one bit of information.

Let κ be the exogenous constant average cost of information acquisition. The overall expected disutility is

$$\sum_x \sum_c h(c|x) g_x(x) L(x,c) + \kappa I(X;C) \tag{1}$$

The citizen's decision problem is to pin down her learning-and-voting choice $h(c|x)$ to minimize (1), subject to a set of constraints that $h(\cdot|x)$ is a probability distribution. This is a standard Kuhn-Tucker problem, and its general solution has been characterized; see Matějka & McKay (2015).

Our case is simple enough for the analytical solution to exist. The following result reports the equilibrium conditions for endogenous marginal and conditional distributions, $\tilde{h}^C(c)$ and $h^C(c|x)$, respectively, where the superscript indicates CA. The proof is in the Appendix. Let $v(x,c) = \exp(-L(x,c)/\kappa)$, let also $w(x) = v(x,0)/v(x,1)$.

Lemma 1. *Optimal marginal distribution $\tilde{h}^C(c)$ is a solution to the equation*

$$\sum_x \frac{v(x,c)}{\sum_{c'} v(x,c') \tilde{h}(c')} g_x(x) = 1$$

for $\tilde{h}(c)$ *(for any* $c = 1, 2$*), as long as this solution is in the interior of* $[0, 1]$*. This equation has an analytical solution,* $\tilde{h}(0) = \frac{1 - w(0) g_x(0) - w(1) g_x(1)}{(1 - w(0))(1 - w(1))}$*. Otherwise,* $\tilde{h}^C(0) = 1$ *or* $\tilde{h}^C(0) = 0$*. Furthermore, the conditional distribution* $h^C(c|x)$ *can be derived from* $\tilde{h}^C(c)$ *as*

$$h^C(c|x) = \frac{v(x, c)\,\tilde{h}^C(c)}{\sum_{c'} v(x, c')\,\tilde{h}^C(c')}$$

In our symmetric case, it is easy to verify that $\tilde{h}^C(0) = \tilde{h}^C(1) = 1/2$ is the solution. Then the conditional distributions are $h^C(c|x)|_{x=c} = (1 + \exp(-1/\kappa))^{-1}$ and $h^C(c|x)|_{x \neq c} = (1 + \exp(1/\kappa))^{-1}$.

I will be interested in normative performance of the assemblies. Firstly, an important criterion is expected loss, secondly, since one may want to see how much information is acquired by the voter, I will also look at mutual information. In case of CA, the expected loss can be calculated to be $EL^C = h(c|x)|_{x \neq c} = (1 + \exp(1/\kappa))^{-1}$, which is a function of κ increasing from 0 to 1/2, initially convex, then switching to concave. Mutual information can be calculated explicitly too, although the formula is tedious.

3 Model of Representative Assembly

The model of the RA is similar to the one in Matějka & Tabellini (2017)—two candidates present their positions in an electoral campaign, the voters acquire information about these positions, and select one candidate. The main difference is that here the voter does not know her own preferences, so she has to acquire information simultaneously about the candidates and about herself—politics and policies.[9]

The collective decision problem is exactly the same as in the CA model. The mechanism through which the decision is made, however, is different. Consider the following game.

Players. There are three agents: two candidates and one citizen-voter. Of course, in reality, there are n citizens voting in this election. However, there is no loss of generality in assuming that there is just one citizen, as, by Assumption 1, each voter accepts that they are pivotal with a constant

[9]There are also other differences. Firstly, the choices of the candidates in Matějka & Tabellini (2017) are policy positions selected with some noise; in my model, the candidate's position is fixed and unknown to the voter, but the candidate's choice of the campaigning intensity may prevent the voters from discovering that type. An alternative interpretation is that candidates in Matějka & Tabellini (2017) have commitment power, while here they do not. Secondly, the policy space is a continuous interval in their model, while it is binary in mine.

probability that can be normalized to one, and by Assumption 2 all voters are identical.

Candidate's preferred policy. Very much like the voter, candidate $k = 1, 2$ has preferences as to which alternative is better. They are described by a random variable Y_k, with the realization $y_k \in \{0, 1\}$. Let $y = (y_1, y_2)$. All three variables X, Y_1, Y_2 come from an exogenous trivariate Bernoulli joint distribution, where the probability of a particular realization is $g(x, y_1, y_2)$. I assume that the associated marginal distributions are always uniform. Beyond that, the key parameters of this probability distribution are the correlation between x and y_k for $k = 1, 2$, denoted ρ_k, and the correlation between Y_1 and Y_2, denoted ρ_0.[10] Correlation ρ_k characterizes the proximity between the voter and candidate k and it will turn out to be a critical exogenous parameter, that characterizes how aligned the candidates and voters are and how severe the agency problem is. Correlation ρ_0 is a measure of proximity between the candidates. For example, $\rho_0 = 1$ represents an ideal coterie, while $\rho_0 = -1$ represents extreme political competition. I follow a plausible assumption that candidates are fully informed agents, or "experts"—they know (x, y), even if voters do not.[11]

Candidate's actions. Once elected, candidate k is programmed to select alternative y_k. In this sense, the candidate is ideological and has no commitment power during the electoral campaign. However, in the run-up to the election, the candidate may try to misrepresent or obfuscate her preferences, in order to appear more acceptable to the voter. In this sense, the candidate is not ideological and just wants to be elected.

Specifically, candidate k generates an electoral campaign message z_k, which is *potentially* observable by the voter. In contrast to z_k, the true preference y_k is not observable by the voter—not even potentially. Therefore, the electoral message z_k may be partially informative about the true y_k, but whatever it is, it still needs to be learned by the voter. The corresponding random variable is Z_k.

Choosing $z_k = y_k$ is costless; however, trying to run a campaign that portrays the candidate as something that she is not, $z_k \neq y_k$, will require

[10]Probability $g(x, y_1, y_2)$ can be expressed fully in terms of seven parameters: three marginal probabilities, three correlations ρ_1, ρ_2, ρ_0, and $\sigma_{x12} = E(X - EX)(Y_1 - EY_1)(Y_2 - EY_2)$. If we assume that marginals are uniform and $\sigma_{x12} = 0$, then ρ_1, ρ_2 and ρ_0 are the only free parameters.

[11]The term expert is not restricted here to refer only to demonstrable scientific facts, such as "vaccines save lives". Someone who advocates in favor of subjective values can also be described as an expert, as long as the CA members would agree with these values after long and careful consideration. For example, many people would agree that "vaccines are unnatural and should be avoided", no matter how long they deliberate.

costly obfuscation effort. For example, the candidate may need to pay off a former love affair, or whitewash some aspects of their past voting record.

Let $\alpha_k(z_k|x, y_k)$ be the probability of z_k. This probability distribution is a strategy of candidate k in the game-theoretic sense. For sake of simplicity, I will parameterize this strategy with a scalar τ_k. Specifically, if the candidate observes that her preferred alternative is the same as that of the voter's, she leaves her electoral message to be the same as the true preferred alternative. In the opposite case, she *attempts* to portray herself as having the same preference as the voter. With probability τ_k she is able to send that message successfully. With the remaining probability $1 - \tau_k$ this attempt is unsuccessful and her message remains equal to her true type. To present this concisely, the following table shows how the probability $\alpha_k(z_k|x, y_k)$ is parameterized by τ_k.

Probability $\alpha_k(z_k|x, y_k)$

	electoral obfuscation $z_k \neq y_k$	electoral truth $z_k = y_k$
no need to obfuscate, $x = y_k$	0	1
obfuscation may be useful, $x \neq y_k$	τ_k	$1 - \tau_k$

In other words, if $\tau_k = 0$ then z_k provides perfect information about y_k, which, together with acquired information about x, enables the voter to learn if the candidate is a good representative. If $\tau_k = 1$ then z_k provides perfect information about x, which alone is useless. Therefore, τ_k represents the level of obfuscation chosen by the candidate in the election. Let $z = (z_1, z_2)$ and $\tau = (\tau_1, \tau_2)$.

Voter's actions. Like in the CA case, I adopt the state-action model, in which the voter chooses directly the probability that candidate $r = 1, 2$ is elected, conditional on learnable x, z. This creates a random variable R; the probability of its realization is denoted $h(r|x, z)$. In terms of notation, k is a generic index of a candidate, while r will tend to denote the identity of the elected representative.

Information acquisition technology. Probability $h(r|x, z)$ is central in this analysis. One needs to clarify which information acquisition technologies are reasonable, that is, whether it makes sense to restrict $h(r|x, z)$ in some way.

It is convenient to view information acquisition as consisting of two parts, *information processing*, and *information transmission*. The processing part is simply preparing a special random variable, some sort of recommendation containing only those aspects of reality that are the most relevant for the voting decision. This recommendation is then transmitted to the voter's brain.

Raw info	Raw info
↓	↓
Info processing	Info transmission
↓	↓
Info transmission	Voter's brain
↓	↓
Voter's brain	Info processing

Figure 1: Information processing outside (left) and inside (right) of the voter

Here is the key statement: costly mutual information, as assumed in this study, penalizes only information transmission, while information processing itself remains costless. Therefore, from the optimization point of view, it makes sense to process information *before* it is transmitted to the voter's brain, to avoid transmitting those aspects that are not relevant for voting. This approach is schematically depicted on the left-hand side of Figure 1.

Leaving $h(r|x,z)$ unconstrained gives the voter the greatest possible level of flexibility in engineering this recommendation.[12] Formally, my benchmark information acquisition technology will conform to the following flexibility assumption.

Assumption 4. The voter can select any probability function $h(r|x,z)$. The minimal quantity of information is $I(X,Z;R)$.

However, this flexibility assumption appears unrealistic at least in some contexts. It requires that there exists some sort of analytical and media infrastructure able to prepare and communicate the recommendation to the individual voter, taking into account both voter's preferences and the alternatives available in the election. That role falls on think-tanks, experts, investigative journalists, pundits, media commentators, etc. I will collectively call these the *commentariat*. Assumption 4 states not only that this commentariat is sophisticated enough to come up with this sort of precise recommendation, but it is also that the voter fully trusts this recommendation.

There is only one case of full flexibility, but there are many ways in which information acquisition may be constrained. Here is one, motivated by friction in information processing. Perhaps it is more realistic to assume that information processing must happen in the voter's brain *after* the information has already been transmitted. This is schematically depicted on the

[12]The concept of flexible information acquisition has been noted in the literature, see for example Denti (2019).

right-hand side of Figure 1. Technically, the voter has to *independently* acquire information about x, z_1 and z_2, and then process the resulting signals into a voting decision. That is, instead of choosing a single conditional distribution $h(r|x,z)$, like in Assumption 4, the voter has to choose three signal distribution functions $h_x(r_x|x)$, $h_1(r_1|z_1)$ and $h_2(r_2|z_2)$, where r_x, r_1 and r_2 are binary signals of x, z_1 and z_2, respectively, and another probability distribution $h_r(r|r_x,r_1,r_2)$ which describes her voting decision r for every combination of received signals. This forms random variables R_x, R_1, R_2 whose realizations are, respectively, r_x, r_1 and r_2.

Assumption 4∗. The voter can select any Bernoulli probability functions $h_x(r_x|x)$, $h_1(r_1|z_1)$, $h_2(r_2|z_2)$ and $h_r(r|r_x,r_1,r_2)$ in order to form

$$h(r|x,z) = \sum_{r_x,r_1,r_2} h_r(r|r_x,r_1,r_2)\, h_x(r_x|x)\, h_1(r_1|z_1)\, h_2(r_2|z_2)$$

The minimal quantity of information is $I(X;R_x) + I(Z_1;R_1) + I(Z_2;R_2)$.

These two assumptions represent alternative views as to how the voter acquires information. The information acquisition technology is far more constrained in Assumption 4∗ than in Assumption 4.[13]

To stress what has already been mentioned—the members of the commentariat are not independent players. The commentariat plays a subservient role to the voters; we should view it as a mechanism shaped by the voters' need to acquire information, subject to assumptions of a particular case at hand, such as Assumption 4 or 4∗.

Election. There are n voters, each voting for k with probability $\tilde{h}(k)$. Assume full correlation, namely, that all voters receive an identical signal from the commentariat so that their votes are all the same.

Assumption 5. Correlation between voters' actions is 1.

Therefore, $\tilde{h}(k)$ is also the probability that k wins the election. Assumption 5 will be relaxed in section 6.

Candidate's preferences. To calculate expected utilities, first define the appropriate probability distributions. Let the joint probability distribution of (x,y,z) be $\beta(x,y,z)$ and let $\gamma(x,z)$ be the associated marginal.[14] Let also $\tilde{h}(r) = \sum_{x,z} h(r|x,z)\,\gamma(x,z)$ be the marginal probability of decision r.

---

[13]Another non-flexible case is when information about electoral campaign can be preprocessed in the form of $h_z(r_z|z_1,z_2)$, and the voter's role is to process independent signals r_x and r_z into a voting decision via $h_r(r|r_x,r_z)$.

[14]$\beta(x,y,z) = \alpha_2(z_2|x,y_2)\,\alpha_1(z_1|x,y_1)\,g(x,y)$ and $\gamma(x,z) = \sum_y \beta(x,y,z)$.

A candidate wants to be elected, receiving payoff 1 when that happens. Ultimately, the candidate k's utility is equal to the probability of winning the election, minus the cost of choosing τ_k,

$$\tilde{h}(k) - \lambda_k \phi(\tau_k). \tag{2}$$

Here, $\lambda_k > 0$ is a parameter measuring how difficult it is to lie, and ϕ is an increasing and convex cost function. In addition to $\phi' > 0$ and $\phi'' > 0$, I assume $\phi(0) = 0$, $\lim_{\tau \to 1} \phi(\tau) = \infty$ and $\phi'(0) = 0$.[15] I will stick to the convention that $\lambda_k = \infty$ implies behavior $\tau_k = 0$.

Voter's preferences. Voter's preferences are described by a disutility function that consists of two elements, exactly like in the CA case. Firstly, she suffers a loss of 1 if the action of the representative r, which is y_r, is different than her own preference characteristic x. That is, the first element is $L(x, y, r) = 1$ if $x \neq y_r$ and zero otherwise. The second element is the cost of acquiring information. This will be captured by mutual information between the decision to elect a candidate, described by a random variable R, and all random variables that can be learned. Under Assumption 4 , the overall expected disutility is

$$\sum_{x,z} \sum_{r} h(r|x, z) \sum_{y} \beta(x, y, z) L(x, y, r) + \kappa I(X, Z; R) \tag{3}$$

Assumption 4$*$ changes only the last term, instead of $I(X, Z; R)$ we have $I(X; R_x) + I(Y_1; R_1) + I(Y_2; R_2)$.

Notice, that this disutility function exploits the assumption that the probability of being pivotal is one.

Solution concept and normative criteria. I consider a simultaneous-move game in which candidate $k = 1, 2$ chooses τ_k and the voter chooses $h(r|x, z)$, under Assumption 4. Under Assumption 4$*$, the voter selects $h_x(r_x|x)$, $h_1(r_1|y_1)$, $h_2(r_2|y_2)$ and $h_r(r|r_x, r_1, r_2)$. In equilibrium, the candidate's choice of τ_k is optimal given the other candidate's choice τ_{-k} and learning-and-voting strategy of the voter, while voter's strategy is optimal given candidates' choices of τ. In short, the solution concept is standard Nash equilibrium. Note that a hypothetical deviation of candidate k to a different τ_k does not affect the strategy of the voter $h(r|x, z)$, but it does affect the probability of being elected $\tilde{h}(r)$.

In order to compare the performance of the assemblies, I will check how often the voter's decision is correct and how much she spends on information acquisition. Let EL^R be the expected loss per voter associated with the incorrect decision in the RA, a counterpart of EL^C in the CA. Alternatively,

[15] In simulations below, I use $\phi(\tau) = \tau^2 / (1 - \tau)$.

it is also the expected probability of the incorrect decision. Let I^R be the mutual information in the RA, a counterpart of I^C in the CA.

There are two remarks regarding welfare calculation. Firstly, the comparison of cost of information acquisition across assemblies is straightforward when there is just one voter. However, when there are many voters, only a tiny fraction of citizens in the CA has to pay a cost of information acquisition, namely only the voter who is selected to be the member of the CA. In the RA, on the other hand, all voters symmetrically pay the same cost of information acquisition, so the true welfare cost of information acquisition is $n\kappa I(X, Z; R)$. By assuming that $n = 1$ in this expression, I underestimate the costs in the RA. Secondly, the population also contains the informed candidates, but their measure is assumed to be zero in welfare calculation.

Solution algorithm

Problem of the voter. The problem of the voter is to find $h(r|x,z)$ in order to minimize objective (3). This is another application of Matějka & McKay (2015). If we define $v(x,z,r)$ to be the exponent of "the expected relative loss from action r, conditional on learnable state (x,z)", or more precisely

$$v(x,z,r) = \exp\left(-\sum_y \frac{\beta(x,y,z)}{\gamma(x,z)} \frac{L(x,y,r)}{\kappa}\right)$$

then the solution of the voter's problem can be characterized in the following way.

Lemma 2. *Optimal marginal distribution $\tilde{h}^R(r)$ is a solution of equation*

$$\sum_{x,z} \frac{v(x,z,r)}{\sum_{r'} v(x,z,r')\,\tilde{h}(r')} \gamma(x,z) = 1$$

for $\tilde{h}(r)$ (for any $r = 1,2$), as long as this solution is in the interior of $[0,1]$; otherwise, $\tilde{h}^R(0) = 1$ or $\tilde{h}^R(0) = 0$. Furthermore, conditional distributions can be derived as

$$h^R(r|x,z) = \frac{v(x,z,r)\,\tilde{h}^R(r)}{\sum_{r'} v(x,z,r')\,\tilde{h}^R(r')}$$

This Lemma gives a best response of the voter to the candidates' choice of (τ_1, τ_2).

Problem of the candidate. The utility of candidate k in (2) can be written more explicitly as

$$\sum_{x,z} h(k|x,z) \sum_y \alpha_k(z_k|x,y_k)\,\alpha_{-k}(z_{-k}|x,y_{-k})\,g(x,y) - \lambda_k c(\tau_k)$$

The first part—the probability of winning the election—is linear in parameter τ_k, so it can be written as $S(\tau_{-k}, h)\,\tau_k$ plus constant, where $S(\tau_{-k}, h)$ stands for the slope coefficient in this term, depending on action of the remaining candidate and the voter, $h(r|x,z)$, $\alpha_{-k}(z_{-k}|x, y_{-k})$ and also on $g(x,y)$. Maximizing the candidate's utility leads to the best response first order condition

$$S(\tau_{-k}, h) - \lambda_k \phi'(\tau_k) = 0. \tag{4}$$

In conclusion, all best response calculations are straightforward (if tedious) and lead to unique solutions. This suggests the following algorithm to find a Nash equilibrium which is employed to find the solutions computationally. For any profile of τ, calculate the best response of the voter, $\tilde{h}(r)$, which we will write as $\tilde{h}^{BR}(\tau)$ to emphasize the dependence on the starting point τ. Then define a function $\Delta(\tau) \in R^2$ as the difference between candidate k's best response and her original action τ_k.

$$\Delta_k(\tau_1, \tau_2) = S\left(\tau_{-k}, \tilde{h}^{BR}(\tau_1, \tau_2)\right) - \lambda_k \phi'(\tau_k)$$

Let τ^R be the solution of $0 = \Delta_k(\tau)$. A pair τ^R and $\tilde{h}^R(r) = \tilde{h}^{BR}\left(\tau^R\right)$ forms a Nash equilibrium, and there are no other Nash equilibria.

4 Benchmark results

The first result pins down the special case when the CA and the RA lead to the same performance. We accept assumptions 1, 2, 3 and 5 throughout this section.

Proposition 1. *Consider any cost of information acquisition $\kappa > 0$. Suppose that all of the following hold: (i) Assumption 4, (ii) $\lambda_1 = \lambda_2 = \infty$, (iii) $\rho_0 = -1$, (iv) $\rho_1 = 0$ and $\rho_2 = 0$. Then $EL^R = EL^C$ and $I^R = I^C$.*

In short, the performance of both assemblies is the same as long as the information acquisition technology is flexible, the cost of electoral obfuscation is prohibitive, the candidates are extremely competitive and no candidate is correlated with the voter.

It is perhaps puzzling that such an equivalence can be achieved. After all, the CA requires learning only about x, while the RA requires learning about both x and z, and that task appears to be harder. It turns out that this intuition is incorrect because, under the conditions of the Proposition, pair (x,z) can be expressed as another uniformly distributed binary random variable. Learning about this recommendation is precisely equivalent to learning about x.

Proof. (Proposition 1) Observe that the prohibitive cost of the electoral obfuscation in condition (ii) implies that the candidates will not attempt to electorally obfuscate, $\tau^R = (0, 0)$, and, consequently, random variables Z and Y will be the same. By condition (iii), for any realization x, there exists a candidate who will implement that x if elected, that is, there is k such that $x = y_k$. These two points together imply that there is a learnable binary recommendation to vote for candidate k if and only if $x = y_k$, which we can denote as R^{rec}. By (i), learning about R^{rec} costs $\kappa I(R^{rec}; R) = \kappa I(X, Z; R)$. By (iv), R^{rec} has a uniform distribution.[16] Therefore, learning about X is as costly as learning about R^{rec} for any κ, and hence the solution in both problems is the same. □

Now, I present some results from comparing the two assemblies for cases that are ruled out by assumptions of Proposition 1. For graphical illustration, these are obtained by solving the models computationally.

Electoral obfuscation. The first case has $\rho_1 = \rho_2 = 0$, $\rho_0 = -1$ and $\lambda_1 = \lambda_2 = 0.3$. Thus, the only difference in comparison to Proposition 1 is that electoral obfuscation is not prohibitively expensive. Figure 2 depicts comparative statics as κ changes. The panel on the left compares the expected loss of two assemblies (the solid curve also shows the expected loss for the RA under parameters of Proposition 1). The panel in the middle makes the same comparison for mutual information. Finally, the panel on the right records the candidates' misinformation strategies τ in the RA. Note that in all cases of equal correlations, $\rho_1 = \rho_2$, the model is symmetric, and so each candidate's ex-ante probability of being selected is $\tilde{h}^R(r) = 0.5$ for any κ.

We see that as κ goes down, the voter exerts more effort to learn the state of nature in both assembly models (mutual information goes up as κ goes down). This coincides with the candidates increasing their effort to obfuscate (τ_k^R goes up as κ goes down), due to the fact that this action is more effective and worth paying for when the voter learns more. In terms of the quality of the decision measured as the expected loss, there are no upshots to using the RA over the CA. In fact, the RA suffers from a double whammy: quality of decision is mechanistically poorer through the noise of the electoral obfuscation, and it is poorer through an equilibrium effect of voters giving up on learning.[17]

[16] Since $\rho_0 = -1$, the remaining correlations must satisfy $\rho_1 + \rho_2 = 0$. Random variable R^{rec} has a uniform distribution if and only if $\rho_1 = \rho_2 = 0$.

[17] This is a good moment to address the question of equilibrium existence. Let us start by observing that equilibrium mutual information does not have to be monotonic in κ (this contrasts with the decision problem). The equilibrium τ^R may become so great as κ

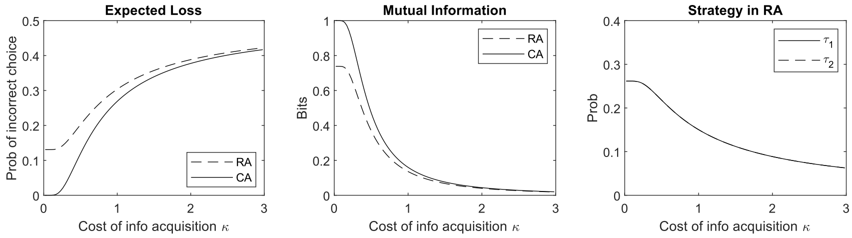

Figure 2: Flexible, $\rho_0 = -1$, $\rho_1 = \rho_2 = 0$ and $\lambda_1 = \lambda_2 = 0.3$.

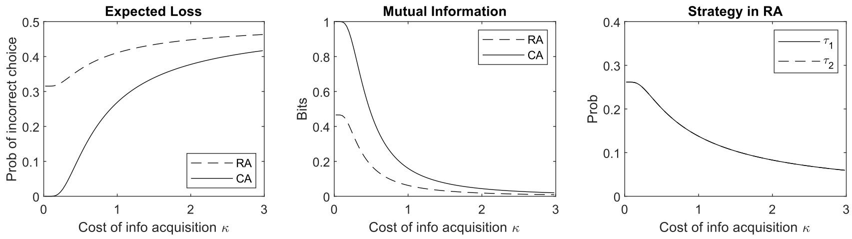

Figure 3: Flexible, $\rho_0 = \rho_1 = \rho_2 = 0$ and $\lambda_1 = \lambda_2 = 0.3$.

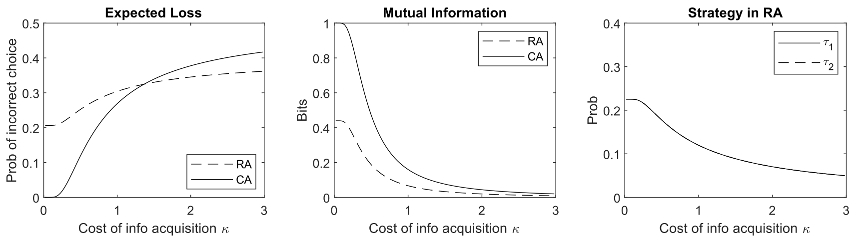

Figure 4: Flexible, $\rho_0 = 0$, $\rho_1 = \rho_2 = 0.2$ and $\lambda_1 = \lambda_2 = 0.3$.

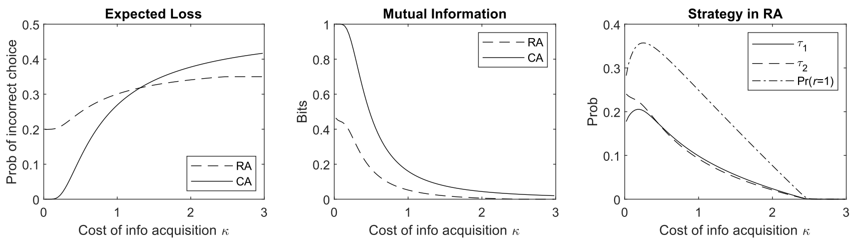

Figure 5: Flexible, $\rho_0 = 0$, $\rho_1 = 0.1$, $\rho_2 = 0.3$ and $\lambda_1 = \lambda_2 = 0.3$.

Coterie. Partial "coterie" is captured by $\rho_0 = 0$, implying that unlucky configurations for the voter, like $x = 0$ and $y_1 = y_2 = 1$, where none of the candidates is a good representative, occur with positive probability. The remaining parameters stay as in the previous case $\rho_1 = \rho_2 = 0$ and (obfuscation) $\lambda_1 = \lambda_2 = 0.3$. The results are depicted in Figure 3 and we see that the CA has a clear advantage over the RA.

Symmetric positive correlation. Now consider a case in which candidates' preferences are somewhat symmetrically correlated with the preferences of the voter, for example, $\rho_1 = \rho_2 = 0.2$. Figure 4 shows that many of the conclusions are similar to the cases discussed above, except one. We see that when κ is high enough, the performance of the RA is better than that of the CA. The reason is simple—because correlations are positive, the voter can always pick one of the candidates and obtain better-than-even odds that this candidate's preferences match her own, even without learning. This opportunity does not exist in the CA.

Asymmetric correlation. Consider a case of asymmetric correlation, $\rho_1 = 0.1$ and $\rho_2 = 0.3$. Now the situation of both candidates is not symmetric ex-ante, and so the equilibrium itself is not symmetric. Figure 5 shows that the voter selects candidate 1 with probability lower than half (this is shown on the right panel as an extra curve denoted as $\Pr(r = 1)$). Moreover, this probability becomes zero for κ high enough. The voter gives up on learning entirely and selects candidate 2 always, relying on the fact that her preferences are more correlated with that candidate's preferences than with those of the other candidate. On the other hand, the candidate's strategies do not diverge from each other that much. The disadvantaged candidate 1 obfuscates only slightly more than candidate 2 for higher κ, but less for low κ.

Constrained information acquisition. Suppose that all parameters are like in Proposition 1 ($\rho_1 = \rho_2 = 0$, $\rho_0 = -1$, and $\lambda_1 = \lambda_2 = \infty$), so, under Assumption 4, the CA is equivalent to the RA. However, let us contemplate an alternative information acquisition technology—the one captured by Assumption 4$*$.

The cost of information acquisition is still linear at rate κ, that is, it is

---

goes down, that it has a detrimental effect on the amount of information acquired by the voter. To illustrate the problem that may arise, assume that $\phi(\tau_k)$ is quadratic, so that $\phi'(\tau_k) = \tau_k$. It is possible to find parameters λ and κ low enough to make $S(\tau_{-k}, h) \geq \lambda$. According to the first-order condition in equation (4), τ_k^R "wants" to become greater than one, and, consequently, mutual information becomes zero. However, it is clear that the upper bound $\tau^R = (1, 1)$ cannot be a Nash equilibrium, and thus no Nash equilibrium in pure strategies exists in this case. The assumption that ϕ diverges to infinity as τ approaches one assures that $\tau_k^R < 1$.

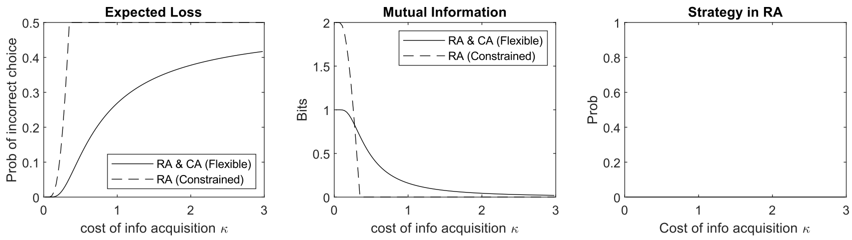

Figure 6: Assumption 4∗, $\rho_0 = -1$, $\rho_1 = \rho_2 = 0$ and $\lambda_1 = \lambda_2 = \infty$.

$\kappa\left(I\left(X; M_x\right) + I\left(Y_1; M_1\right) + I\left(Y_2; M_2\right)\right)$.

Figure 6 compares the result of Proposition 1 with the result of this alternative formulation. We see that the expected losses from both models converge to zero when κ converges to zero. But as κ increases, the expected loss explodes in the case of Assumption 4∗ quite more dramatically than in the optimal information acquisition version of the model. As κ approaches zero, the total mutual information is twice as high as in the case of full flexibility. That is, the voter is satisfied with using two bits of information, as there are two binary independent random variables x and y_1 to learn about (where $y_2 = 1 - y_1$).

It has been hypothesized that the CA may play a role even if it is not responsible for the final decision (Warren & Gastil 2015). Assemblies may be organized as deliberative polling events (Fishkin 2011), whose main objective is to provide information to citizens who will take a consequential decision by voting in the general election. Such *advisory* CA may complement media, experts, or pundits who sometimes distort the truth and who are not always trusted by the general public. In the context of the analysis of this section, this is equivalent to removing the constraints like the ones in Assumption 4∗, so that we are closer to the situation in Assumption 4. On the other hand, there is no reason to believe that this advisory CA can easily overcome other distortions, such as electoral obfuscation or coterie. They may investigate the candidates more, but this also means that the candidates could react by spending more on obfuscation. Either way, advisory CA is not as efficient as the pure CA studied in this section.[18]

[18] However, an advisory CA could replicate the pure CA if it is followed by a final decision referendum, rather than by a general election of representatives who make a decision. But then this configuration is almost indistinguishable from a pure CA.

5 Hybrid model—Citizen Assembly with experts

The most important parameters of the model are correlations between the voter and the candidates, ρ_1 and ρ_2. Each can be interpreted as a degree to which the candidate can be treated as a trustworthy expert. The voter can simply lean on the opinion of the candidate, in a similar way we rely on doctors, architects, or car mechanics when we follow their recommendations about the correct course of action. Since the pure CA model does not have such experts, it is not able to benefit from their decisions.

Suppose, however, that the CA can invite outsiders to the proceedings and follow their advice. Assume that the agents who played the role of the candidates in the pure RA are now repurposed as "experts". In other words, there is the CA element—a direct choice between status quo and reform—and the RA element—an option to entrust one of two experts with the decision. With this slight modification, the CA model might be able to replicate all the benefits of RA, and sometimes perform strictly better.

Model

Experts are equivalent to candidates in the sense of having the same ρ_k. As far as experts' preferences are concerned, it is not clear whether they should be the same as those of the political candidates in the previous section. Therefore, I will consider two cases. In the first one, a sincere expert just mechanically offers her services; this is captured by $z_k = y_k$, or simply $\lambda_k = \infty$. In the second formulation, experts are the same as the candidates in the pure RA model in the sense of deriving utility from winning the argument and being able to obfuscate their message to improve their chances of winning this argument; their utility is like in (2) with some $\lambda_k > 0$. For instance, they may have career concerns that are served by winning the argument.

As far as the voter's decision is concerned, it comes from the set of four possibilities, $d \in \{0_c, 1_c, 1_r, 2_r\}$. Here, 0_c and 1_c stand for selecting, respectively, the status quo and reform, while 1_r and 2_r mean following the recommendation of, respectively, expert 1 and expert 2. This gives rise to a random variable D with the realization d; voter's mutual information is $I(X, Z; D)$. The overall disutility is constructed in the straightforward way.

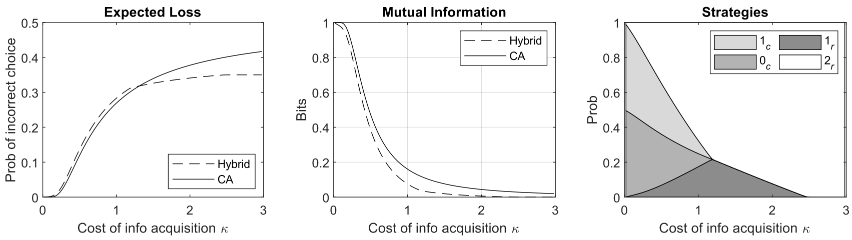

Figure 7: Hybrid, flexible, $\rho_0 = 0$, $\rho_1 = 0.1$, $\rho_2 = 0.3$ and $\lambda_1 = \lambda_2 = 0.3$.

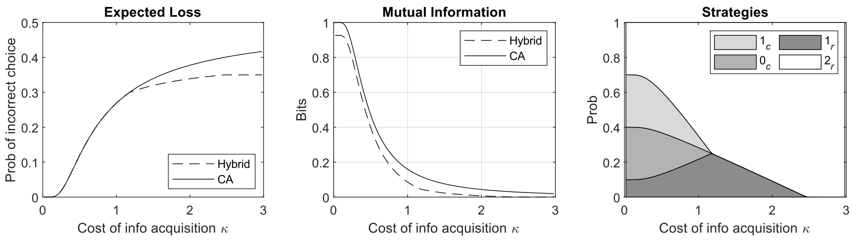

Figure 8: Hybrid, flexible, $\rho_0 = 0$, $\rho_1 = 0.1$, $\rho_2 = 0.3$ and $\lambda_1 = \lambda_2 = \infty$.

Results

All simulations assume flexible information acquisition and asymmetric strictly positive correlations ρ_1 and ρ_2. Figure 7 shows obfuscating experts, $\lambda_k = 0.3$, while Figure 8 presents the case of non-lying experts, $\lambda_k = \infty$.[19]

The key take-home message is not that surprising—the hybrid model has the advantages of both the CA and RA. Specifically, there are three cases.

For low costs of information acquisition, the voter takes any of the four decisions with positive probability. The right-hand side panel in the Figures shows this when κ is to the left of the point where all the gray areas meet. The lower the cost, the more the voter relies on direct decisions, 0_c and 1_c. In fact, in the case when experts have incentives to obfuscate, as in Figure 7, the probability that experts are selected converges to zero as $\kappa \to 0$. However, if experts are fully sincere then they are selected with probability bounded away from zero even for arbitrarily small costs of information acquisition. This is interesting because it indicates a non-continuity of comparative statics around point $\kappa = 0$; note that when the information is costless, the voter would select an expert with probability zero.

[19]The right panels of these Figures show strategies of the voters—the vertical distance is the probability of each of four possible decisions. The light gray areas show the probabilities of direct decisions 0_c and 1_c, while the white and dark areas show the probabilities of indirect decisions 1_r and 2_r.

For higher costs of information acquisition, the voter ceases to take a direct decision, $\Pr(0_c) = \Pr(1_c) = 0$, but acquires information in an effort to decide which expert, 1_r or 2_r, to select, and selects both of them with positive probability.

Finally, at the other end of the spectrum, for costs of information acquisition high enough (for the generic case of non-equal correlation coefficients $\rho_1 \neq \rho_2$), the voter does not acquire any information but relies on the expert with the highest correlation coefficient.

What is important is this: when the cost of information acquisition is low, the hybrid model is *strictly* more efficient than either the pure RA or the pure CA. This follows directly from the observation that all four possible decisions are used by the CA with positive probability.[20]

6 Other extensions

6.1 Independent voters

In the analysis of the pure RA model above, the signals that the voters were receiving from the commentariat were perfectly correlated, leading to perfectly correlated votes. Instead, suppose now that they are independent.

If voters are independent, then the number of votes in favor of candidate k would follow a binomial distribution with parameters n and probability of success in a single draw $\tilde{h} = \tilde{h}(k)$. Even a slight asymmetry in the form of $\tilde{h} \neq 1/2$ would translate into a dramatic asymmetry in the overall probability of winning, as it does in binomial distributions. I am interested in the implication of this assumption on the equilibrium electoral obfuscation.

Since we consider the case of large n, we can take advantage of the normal approximation of the binomial distribution. Let $v = 0, 1, ..., n$ be the number of votes in favor of candidate k, then the approximate distribution of v is $N\left(n\tilde{h}, n\tilde{h}\left(1-\tilde{h}\right)\right)$. Candidate k wins if $v > n/2$, and the probability of this event is

$$\Pr(k \text{ wins}) = 1 - \Phi\left(\sqrt{n}\frac{\frac{1}{2} - \tilde{h}}{\sqrt{\tilde{h}\left(1-\tilde{h}\right)}}\right)$$

where Φ is standard normal c.d.f.

The pure RA model is exactly the same as in section 3 with one exception. The utility of the candidate in equation (2) is replaced by $\Pr(k \text{ wins}) -$

[20] The voter wants to minimize a combination of the expected loss and mutual information. Therefore, the expected loss on the left panels of the Figures may be slightly higher than a pure assembly type, and yet the overall disutility is still lower.

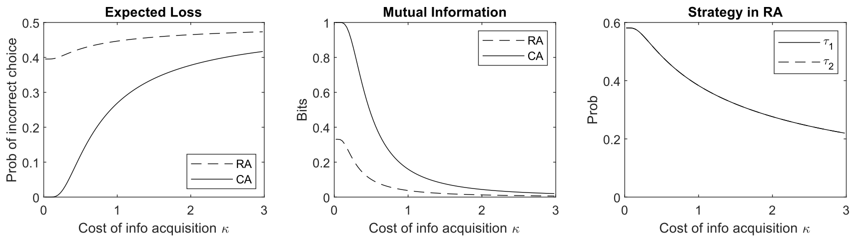

Figure 9: Case of independent voters, $n = 50$, $\rho_0 = \rho_1 = \rho_2 = 0$ and $\lambda_1 = \lambda_2 = 0.3$.

$\lambda_k \phi(\tau_k)$, and the candidate's first order condition (4) with respect to τ_k becomes now

$$\frac{d \Pr(k \text{ wins})}{d\tilde{h}} S(\tau_{-k}, h) - \lambda_k \phi'(\tau_k) = 0.$$

Figure 9 depicts the case of $n = 50$ independent voters; other parameters are like in Figure 3. As compared to the correlated case, gaining a slight advantage in this scenario is more profitable for each candidate, and this triggers more obfuscation in equilibrium, which in turn has a detrimental effect on learning by the voters. This worse learning outcome in a large population has nothing to do with endogenous probability of being pivotal and the paradox of voting, which are assumed away in this analysis.

6.2 Many issues

In any society, there are many collective decisions to be taken at the same time. In this subsection, I assume that there are M issues which are drawn *independently* from the distribution g. Issue m is characterized by a triple (x^m, y_1^m, y_2^m)

The critical question is how different types of assembly cope with this multiplicity. I adopt the following institutional assumption.

Assumption 6. Representative Assembly: there is one general election resuling in an assembly deciding about all M issues. Citizen Assembly: there are M independent assemblies drawn from the population of citizens (assuming $M \leq N$), each deciding about one issue.

Note that different arrangements, while conceptually possible, seem completely implausible or absurdly impractical. For example, one could contemplate separate general elections for different issues. Similarly, one could have a single draw into one CA that would subsequently address all issues.

The claim of this section is that with many issues the information acquisition problem for the RA is even more acute than in the CA.

To see this, notice that under Assumption 6, the CA system is perfectly scalable with respect to the number of issues. The multi-issue RA system, on the other hand, suffers from a novel problem that has not occurred in the analysis above. For example, if the cost of information acquisition is small, then the voter can pick the correct candidate with high probability in a single issue case. In a multi-issue case, however, one candidate has to be selected for all issues, by Assumption 6. The voter will pick the candidate whose preferences match her preferences best, but even that candidate will have mismatched preferences in some issues. The result is that the expected loss will be worse.

Notice that this is a different problem than the one caused by many issues making the information acquisition problem more difficult. For example, under the full flexibility Assumption 4, the information acquisition problem is not more acute, as the commentariat's "voting recommendation" can still be simple binary.

7 Conclusions

The CA and RA are different in so many ways that comparing them is not a straightforward enterprise. Confronted with such a task, the first step is to find a configuration of parameters that makes the two assemblies equivalent, and then change only one aspect at a time, ceteris paribus, to see its effect.

The key challenge of any collective choice method is to utilize the expertise that hides somewhere in the society, without falling victim to a potential agency problem. The RA system attempts to take advantage of the information that candidates possess, but—before handing them power—it asks the voters to learn which of these candidates is more suitable. I argue that learning is not a friction-less process, and I show that this can be efficient only if the agency problem is not too severe, namely when the candidates' and the voters' preferences are ex-ante sufficiently correlated. If not, then the do-it-yourself approach of the pure CA is better.

The main practical question is how to construct a hybrid system that takes advantage of the strength of both methods. For example, Warren & Gastil (2015) imagine the CA as a trusted information-gathering institution without decision-making powers; voters could use its recommendations to decide about votes in otherwise normal general election. This could be advantageous in moving us from the reality of Assumption 4$*$ to Assumption 4. However, I would argue that this approach is unlikely to solve other problems, such as, for example, electoral obfuscation.

The hybrid model proposed here assumes that the CA is not merely advisory but has real decision-making powers. According to this model, the informed agents (i.e. those who were candidates in the pure RA system) are invited to the CA in order to provide their expert testimony. For low information acquisition cost, this hybrid system is strictly better than the pure CA or RA. More intricate types of hybrid models can be considered, and this is left for future research.

There are other arguments in favor of the CA. If there are independent multiple issues then it would be easier to fine-tune the decision if we had a series of independent assemblies rather than a single big one; this is easier to achieve with CAs rather than RAs.

Having many voters can affect the results in a number of ways. In this paper, I consider a RA with voters voting independently. I computationally show that, in comparison to correlated voting, there are greater incentives for the candidates to obfuscate, which has an additional detrimental effect on learning by the voters. This has nothing to do with the effect that the endogenous probability of being pivotal would have. However, the paradox of voting is bound to be at play in the large-scale general elections, while it is likely to be insignificant in a small-group CA.

Another parameter of the model that I kept constant in my analysis is the cost of information acquisition, κ. However, this too is easier to manipulate in the CA than in the RA. One can control the opportunity cost of time of citizens who are locked for a few days in a CA. One could restrict access to other distracting activities such as web-browsing or TV entertainment and thus make κ very low, while such a control does not exist in the general elections to the RA.[21]

Finally, I omitted the case in which voters are heterogeneous. This could be preference heterogeneity as in, for example, the Hotelling tradition, or different access to information as in, for example, Bardhi & Bobkova (2020). This aspect is probably the most interesting direction for future research.

References

Bardhi, A. & Bobkova, N. (2020), 'Deliberative minipublics', *Presentation at the SITE Stanford, 31 July* .

Ben-Yashar, R. & Nitzan, S. (2001), 'The invalidity of the condorcet jury

[21]Lower κ would increase information acquisition, and therefore its cost, so one could argue that the overall welfare effect is unclear. However, the cost of information acquisition in the CA is socially negligible, as only a tiny fraction of all citizens gets to sit in the assembly and pay that cost. In other words, in the CA, we should be looking at the social expected loss, and ignore the social cost of information acquisition.

theorem under endogenous decisional skills', *Economics of Governance* **2**(3), 243–249.

Bhattacharya, S., Duffy, J. & Kim, S. (2017), 'Voting with endogenous information acquisition: Experimental evidence', *Games and Economic Behavior* **102**, 316–338.

Denti, T. (2019), 'Unrestricted information acquisition', *Unpublished manuscript Cornell University* .

Downs, A. (1957), *An economic theory of democracy*, New York, Harper.

Fishkin, J. S. (2011), *When the people speak: Deliberative democracy and public consultation*, Oxford University Press.

Gastil, J. & Wright, E. O., eds (2019), *Legislature by lot: Transformative designs for deliberative governance.*, Verso Books.

Gerardi, D. & Yariv, L. (2008), 'Information acquisition in committees', *Games and Economic Behavior* **62**(2), 436–459.

Koriyama, Y. & Szentes, B. (2009), 'A resurrection of the condorcet jury theorem', *Theoretical Economics* **4**(2), 227–252.

Manin, B. (1997), *The Principles of Representative Government*, Cambridge University Pr.

Martinelli, C. (2007), 'Rational ignorance and voting behavior', *International Journal of Game Theory* **35**(3), 315–335.

Matějka, F. & McKay, A. (2015), 'Rational inattention to discrete choices: A new foundation for the multinomial logit model', *American Economic Review* **105**(1), 272–98.

Matějka, F. & Tabellini, G. (2017), 'Electoral competition with rationally inattentive voters', *Available at SSRN 3070204* .

Mukhopadhaya, K. (2003), 'Jury size and the free rider problem', *Journal of Law, Economics, and Organization* **19**(1), 24–44.

Oliveros, S. (2013), 'Abstention, ideology and information acquisition', *Journal of Economic Theory* **148**(3), 871–902.

Paulis, E., Pilet, J.-B., Panel, S., Vittori, D. & Close, C. (2020), 'The politicize dataset: an inventory of deliberative mini-publics (dmps) in europe', *European Political Science* pp. 1–22.

Persico, N. (2004), 'Committee design with endogenous information', *The Review of Economic Studies* **71**(1), 165–191.

Smith, G. (2005), *Beyond the ballot: 57 democratic innovations from around the world*, Power Inquiry.

Smith, G. (2009), *Democratic innovations: Designing institutions for citizen participation*, Cambridge University Press.

Warren, M. E. & Gastil, J. (2015), 'Can deliberative minipublics address the cognitive challenges of democratic citizenship?', *The Journal of Politics* **77**(2), 562–574.

A Appendix: Proofs

A.1 Proof of Lemma 2

This is similar to Matějka & McKay (2015).

Let $\xi(x,z)$ be the Lagrangian multiplier on constraint $\sum_r h(r|x,z) = 1$. The Lagrangian can be written as

$$J = \sum_{x,z,r} h(r|x,z) \sum_y \beta(x,y,z) L(x,y,r) + \\ + \kappa I(x,z;r) + \sum_{x,z} \xi(x,z) \sum_r h(r|x,z)$$

Its derivative w.r.t. voter's choice variable $h(r|x,z)$ for any (r,x,z) is

$$\frac{\partial J}{\partial h(r|x,z)} = \sum_y \beta(x,y,z) L(x,y,r) + \kappa\gamma(x,z) \ln \frac{h(r|x,z)}{\tilde{h}(r)} + \xi(x,z)$$

Equalizing to zero and rearranging gives

$$-\frac{\sum_y \beta(x,y,z) L(x,y,r)}{\kappa\gamma(x,z)} = \ln \frac{h(r|x,z)}{\tilde{h}(r)} + \frac{\xi(x,z)}{\kappa\gamma(x,z)}$$

Define $\ln \mu(x,z) = \xi(x,z) / (\kappa\gamma(x,z))$ and

$$v(x,z,r) = \exp\left(-\sum_y \frac{\beta(x,y,z)}{\gamma(x,z)} \frac{L(x,y,r)}{\kappa}\right)$$

Then f.o.c. becomes

$$v\left(x,z,r\right)\tilde{h}\left(r\right)=h\left(r|x,z\right)\mu\left(x,z\right) \tag{5}$$

Sum over r and use the constraint $\sum_r h\left(r|x,z\right)=1$ to obtain

$$\sum_r v\left(x,z,r\right)\tilde{h}\left(r\right)=\mu\left(x,z\right)$$

Plug back to (5) to eliminate $\mu\left(x,z\right)$

$$\frac{v\left(x,z,r\right)\tilde{h}\left(r\right)}{\sum_{r'} v\left(x,z,r'\right)\tilde{h}\left(r'\right)}=h\left(r|x,z\right)$$

Multiply by $\gamma\left(x,z\right)$ and sum over x,z. Then, we obtain for all r

$$\sum_{x,z}\frac{v\left(x,z,r\right)}{\sum_{r'}^{2} v\left(x,z,r'\right)\tilde{h}\left(r'\right)}\gamma\left(x,z\right)=1$$

So far, this calculation ignored the constraint that $0\leq\tilde{h}\left(r'\right)\leq 1$, but it is trivial to impose it on the candidate solution $\tilde{h}\left(r\right)$.

A.2 Appendix: constrained model

This section investigates a model under Assumption 4$*$, and in which the voter faces a decision problem (thus, not a game, assume $\lambda_1=\lambda_2=\infty$ or simply that $Z=Y$). We imagine that the voter establishes four conditional distributions $h_r\left(r|r_x,r_1,r_2\right)$, $h_x\left(r_x|x\right)$, $h_1\left(r_1|y_1\right)$ and $h_2\left(r_2|y_2\right)$ with interpretation that the voter summarizes x in random variable r_x, y_k in r_k, and then summarizes r_x,r_1,r_2 in voting decision r. All these probabilities relate to our familiar $h\left(r|x,y\right)$ via

$$h\left(r|x,y\right)=\sum_{r_x,r_1,r_2} h_r\left(r|r_x,r_1,r_2\right)h_1\left(r_1|y_1\right)h_2\left(r_2|y_2\right)h_x\left(r_x|x\right)$$

However, the point of this section is that only those $h\left(\cdot\right)$ that can be decomposed into $h_r\left(\cdot\right)$, $h_1\left(\cdot\right)$,$h_2\left(\cdot\right)$ and $h_x\left(\cdot\right)$ are allowed.

The cost of information acquisition remains linear in mutual information with a coefficient κ.

The overall disutility function is

$$D=\sum_{x,y,r,r_x,r_1,r_2}\begin{pmatrix} h_r\left(r|r_x,r_1,r_2\right)\times \\ h_1\left(r_1|y_1\right)\times \\ h_2\left(r_2|y_2\right)\times \\ h_x\left(r_x|x\right)\end{pmatrix} g\left(x,y\right)L\left(x,y,r\right)+\kappa\begin{pmatrix} I\left(X;R_x\right)+ \\ I\left(Y_1;R_1\right)+ \\ I\left(Y_2;R_2\right)\end{pmatrix}$$

The endogenous terms $h_r(\cdot)$, $h_1(\cdot)$,$h_2(\cdot)$ and $h_x(\cdot)$ are conditional distribution functions so let $\xi_x(x)$, $\xi_1(y_1)$, $\xi_2(y_2)$ and $\xi(r_x, r_1, r_2)$ be the associated Lagrange multipliers. The Lagrangian is

$$D + \left(\begin{array}{c} \sum_x \xi_x(x) \sum_{r_x} h_x(r_x|x) + \\ \sum_{y_1} \xi_1(y_1) \sum_{r_1} h_1(r_1|y_1) + \\ \sum_{y_2} \xi_2(y_2) \sum_{r_2} h_2(r_2|y_2) \end{array} \right) + \sum_{r_x, r_1, r_2} \xi(r_x, r_1, r_2) \sum_r h_r(r|r_x, r_1, r_2)$$

The solution proceeds as follows.

Step 1. Optimize with resp to $h_x(r_x|x)$ for every (r_x, x). Foc is

$$-\sum_{y,r,r_1,r_2} \left(\begin{array}{c} h_r(r|r_x, r_1, r_2) \times \\ h_1(r_1|y_1) \times \\ h_2(r_2|y_2) \times \end{array} \right) g(x,y) L(x,y,r) = \kappa g_x(x) \ln\left(\frac{h_x(r_x|x)}{\tilde{h}_x(r_x)}\right) + \xi_x(x)$$

Following the usual steps, we can define

$$v_x(x, r_x) = \exp \sum_{y,r,r_1,r_2} \left(\begin{array}{c} h_r(r|r_x, r_1, r_2) \times \\ h_1(r_1|y_1) \times \\ h_2(r_2|y_2) \times \end{array} \right) \frac{g(x,y)}{g_x(x)} \frac{-L(x,y,r)}{\kappa}$$

and characterize the foc as

$$\sum_x \frac{v_x(x, r_x)}{\sum_{r'_x} v_x(x, r'_x) \tilde{h}_x(r'_x)} g_x(x) = 1$$

Of course, this time we cannot solve explicitly for endogenous $\tilde{h}_x(r_x)$ because $v(x, r'_x)$ contains endogenous $h_r(\cdot)$, $h_1(\cdot)$ and $h_2(\cdot)$.

Step 2. Following similar steps, we can characterize optimal $\tilde{h}_1(r_1)$, which will depend on $h_r(\cdot)$, $h_x(\cdot)$ and $h_2(\cdot)$.

Step 3. Following similar steps, we can characterize optimal $\tilde{h}_2(r_2)$.

Step 4. Now find optimal $h_r(r|r_x, r_1, r_2)$ for every r, r_x, r_1, r_2. Our Lagrangian is linear in $h_r(\cdot)$ and, therefore, the optimal solution is 0 or 1. Thus, if, for a particular r_x, r_1, r_2, we have

$$\sum_{x,y} \left(\begin{array}{c} h_x(r_x|x) \times \\ h_1(r_1|y_1) \times \\ h_2(r_2|y_2) \end{array} \right) g(x,y) L(x,y,1) > \sum_{x,y} \left(\begin{array}{c} h_x(r_x|x) \times \\ h_1(r_1|y_1) \times \\ h_2(r_2|y_2) \end{array} \right) g(x,y) L(x,y,2)$$

then the voter should vote for the second candidate, $h(1|r_x, r_1, r_2) = 0$.

To summarize, we can imagine the following algorithm starting from initial $h_x(r_x|x)$, $h_1(r_1|y_1)$ and $h_2(r_2|y_2)$.

1. Calculate $h(1|r_x, r_1, r_2)$ for these functions using step 4.

2. For initial conditional distributions $h_x(\cdot)$, $h_1(\cdot)$ and $h_2(\cdot)$ find the next iteration of marginal distributions $\tilde{h}_x(\cdot)$ $\tilde{h}_1(\cdot)$ and $\tilde{h}_2(\cdot)$ using steps 1, 2 and 3. Then recover the conditional distributions. Denote this next iteration as Ψ_x, Ψ_1 and Ψ_2.

3. Iterate. We are looking for (h_x, h_1, h_2) that solves system $0 = \Psi_x(h_x, h_1, h_2) - h_x$, $0 = \Psi_1(h_x, h_1, h_2) - h_1$ and $0 = \Psi_2(h_x, h_1, h_2) - h_2$.