A Recursive Measure of Voting Power with Partial Decisiveness or Efficacy

Arash Abizadeh

[Forthcoming in *Journal of Politics*]

Abstract: The current literature standardly conceives of voting power in terms of decisiveness: the ability to change the voting outcome by unilaterally changing one's vote. I argue that this classic conception of voting power, which fails to account for *partial* decisiveness or efficacy, produces erroneous results because it saddles the concept of voting power with implausible microfoundations. This failure in the measure of voting power in turn reflects a philosophical mistake about the concept of social power in general: a failure to recognize that an agent can exercise individual social power with others' assistance, in virtue of the group's collective power, sometimes even when she could not unilaterally scuttle the group's collective power. I therefore develop a conception of efficacy that admits of degrees and defend a Recursive Measure of voting power that takes partial efficacy into account.

Keywords: social power; voting power; decisiveness; collective power; power-with

Funding: Support for this research was provided by Social Sciences and Humanities Research Council of Canada, Insight Grant 453-2018-0386

The distribution of power lies at the heart of the study of politics: in empirical studies, we want to know who has power to determine political outcomes; in normative studies of institutional and especially constitutional design, we want to know the extent to which different decision-making rules conform to normative standards that require maximizing or equitably distributing decision-making power. Given the centrality of voting to political decision-making, a satisfactory measure of voting power is indispensable.

In the current literature, voting power is standardly conceived in terms of *decisiveness*, i.e., a voter's ability to change the outcome by unilaterally changing her vote. (For example, under majority-rule voting with five voters, in a *division* or complete vote configuration with three YES-votes and two NO-votes, each YES-voter is decisive, because if she were to have voted differently the outcome would have been different.) My critical thesis is that the classic reduction of voting power to the potential to be decisive in this sense—*fully* decisive or efficacious—saddles the concept of voting power with implausible microfoundations: it presupposes a mistaken analysis of the power that individuals efficaciously exercise in particular divisions. This reduction falsely depicts causally *overdetermined* outcomes, for which no one is (fully) decisive, as not resulting from any voter's efficacious power. Moreover, these implausible microfoundations lead to evidently mistaken outputs—erroneous measures of voting power—once the distribution of preferences (and strategic interaction) is taken into account. My constructive thesis is that recognizing partial decisiveness—or, to use a more intuitive label, *partial efficacy*—remedies both problems. I therefore defend a conception of efficacy or decisiveness admitting of degrees and, accordingly, a recursive measure of voting power taking partial efficacy into account. Whereas the classic measure is typically interpreted as representing voters' *probability* of being decisive, the Recursive Measure represents their *expected efficacy*.

The failure to recognize partial efficacy in measuring voting power reflects a philosophical mistake about social power more generally: a failure to recognize that one can exercise individual agential power in effecting outcomes *with* others' assistance, in virtue of the group's collective power, sometimes even when one could not unilaterally scuttle the group's collective power (Abizadeh 2021b). An agent efficaciously exercises power not only when her action is necessary for the outcome, but also when it is a necessary element of a set of conditions sufficient for the outcome, which set is a subset of the actual set of conditions effecting it (assuming the causal process, of which it is an element, is not pre-empted). A measure taking this into account is of broad significance. First, the Recursive Measure resolves a number of puzzles facing the classical model. Second, some of these puzzles, such as how to explain the unequal power suffered by members of persistent minorities under majoritarian decision-making, are of great empirical and normative significance. Finally, given its formal-institutional setting, voting power is an especially tractable instance of social power: clarifying its nature helps clarify the concept of social power more generally, and provides an analytically clean benchmark for studying it in other, messier settings.

The argument proceeds as follows. First, I introduce the classic Penrose-Banzhaf measure of a priori power, generalize it to cover a posteriori power, and make explicit the classic model's three underlying assumptions. Next, I show why one of these assumptions—which reduces efficacy to full decisiveness—presupposes a mistaken analysis of causally overdetermined divisions, argue that the assumption cannot be salvaged by explaining such divisions in terms of a group's collective power, and then show that, because of this assumption, the classic model produces mistaken measures of a posteriori power. Then, I argue that causal overdetermination is explicable via the notion of partial efficacy, critically examine previous attempts to articulate such a notion, and, with corrections in hand, defend and illustrate a recursive measure of voting power taking partial efficacy into account. Finally, I provide an intuitive interpretation of this measure, demonstrate its power to resolve puzzles plaguing the classic model, and conclude with the significance of the Recursive Measure for empirical and normative analysis.

A Priori Voting Power

The standard measures of voting power, the *Penrose-Banzhaf Measure* (*PB*) and *Shapley-Shubik Index* (*SS*), conceive of voting power in terms of agents' potential causal contribution to or efficacy over outcomes: both are examples of what I call *Efficacy Measures* (*EM*) of voting power. Both were also initially conceived of as measures of so-called *a priori* voting power, which refers to the voting power one has solely in virtue of the formal-institutional structure of the decision-making procedure once an agenda-set of alternatives is given, i.e., the sets of alternative outcomes, actors, their action profiles, and the collective-decision function that maps combinations of actions onto outcomes (Felsenthal and Machover 1998, 2003, 2004). *A posteriori* voting power, by contrast, includes the power one has in virtue of other social structures, such as class, ethnicity, or gender, and the ensuing distribution of preferences (and strategic incentives). I here focus on *PB* for two reasons: *SS*'s suitability as a measure of a priori voting power has been put into doubt; and, more importantly, interpreted in probabilistic terms, *SS* amounts to a special a posteriori case of a generalization of *PB*. (See Appendix on *SS*.)

Any Efficacy Measure of a voter v_i's voting power can be represented as:

[EM] $$EM_i = \sum_{d_j \in D} e_i(d_j) \cdot w(d_j)$$

where $D=\{d_1, d_2 \ldots d_{|D|}\}$ is the set of all $|D|$ logically possible divisions of the set $V=\{v_1, v_2 \ldots v_{|V|}\}$ of voters; $e_i(d_j)$ is v_i's *division efficacy score* in d_j; and $w(d_j)$ is the *division weight* assigned to any $d_j \in D$. What is distinct about each measure is its specification of $e_i(d_j)$ and $w(d_j)$. Now, *PB* equates a priori voting power with the proportion of logically possible divisions, within a voting structure, in which an agent is decisive or has full causal efficacy, i.e., in which, holding everyone else's votes fixed, she could have altered the outcome by changing her vote (Felsenthal and Machover 1998, 2004).[1] Its division efficacy scores and division weights are therefore defined as equal to:

[E.1] $$e_i(d_j) = \begin{cases} 1 & \text{if } v_i \text{ is decisive in division } d_j \\ 0 & \text{otherwise} \end{cases}$$

[W.1] $$w(d_j) = \frac{1}{|D|}$$

To illustrate, consider three voters using majority-rule voting (M3) versus unanimity-rule voting (U3) who face two alternatives (YES or NO) and can vote YES or NO. We represent each voting structure in curly brackets by the quota for a YES-outcome, followed after the semicolon by the weight of each voter's vote. In Tables 1 and 2, we list the vote $a_i(d_j)$ each voter v_i casts in each division d_j, paired with her efficacy score $e_i(d_j)$, underlining decisive YES-votes and bolding decisive NO-votes. $\sum PB_i$ represents the *overall a priori voting power* in a given voting structure.

M3 and U3 illustrate three important facts about *PB* and the concept of voting power it presupposes. First, the measure assumes that more than one voter can be decisive in effecting an outcome. This is why the values of $e_i(d_j)$ in a single division can sum up to more than 1. The value of $e_i(d_j)$ does not represent a voter's relative share of efficaciously exercised voting power: any action necessary for the outcome under the circumstances is decisive. Call this the *non-constant-sum-efficacy* assumption.[2] Second, overall voting power can vary from structure to structure: because each voter has greater voting power under M3 than under

Table 1: Majority-rule Voting M3={2;1,1,1}

Divisions *D*	$a_i(d_j)$, $e_i(d_j)$			Outcome
(|D|=2^3=8)	v_1	v_2	v_3	
d_1	y,0	y,0	y,0	y
d_2	n,0	y,1	y,1	y
d_3	y,1	n,0	y,1	y
d_4	**n**,1	**n**,1	y,0	n
d_5	y,1	y,1	n,0	y
d_6	**n**,1	y,0	**n**,1	n
d_7	y,0	**n**,1	**n**,1	n
d_8	n,0	n,0	n,0	n
PB_i	$\frac{4}{8}$ = ½	$\frac{4}{8}$ = ½	$\frac{4}{8}$ = ½	$\sum PB_i$=1½

Table 2: Unanimity-rule Voting U3={3;1,1,1}

Divisions *D*	$a_i(d_j)$, $e_i(d_j)$			Outcome
(|D|=2^3=8)	v_1	v_2	v_3	
d_1	y,1	y,1	y,1	y
d_2	**n**,1	y,0	y,0	n
d_3	y,0	**n**,1	y,0	n
d_4	n,0	n,0	y,0	n
d_5	y,0	y,0	**n**,1	n
d_6	n,0	y,0	n,0	n
d_7	y,0	n,0	n,0	n
d_8	n,0	n,0	n,0	n
PB_i	$\frac{2}{8}$ = ¼	$\frac{2}{8}$ = ¼	$\frac{2}{8}$ = ¼	$\sum PB_i$=¾

U3, $\sum PB_i$ under M3 is greater than under U3. Overall voting power is not constant sum. Third, although an individual's voting power ranges from 0 (a *dummy* not decisive in any division) to 1 (a *dictator* decisive in every division), $\sum PB_i$ can be greater than 1.

Interpretation and Generalization: A Posteriori Voting Power

The standard interpretation of *PB* is that it represents the probability a voter will be decisive in a given voting structure under the following assumptions: (a) *equiprobable voting* (the probability a voter votes for an alternative is equal to the probability she votes for any other) and (b) *voting independence* (votes are not correlated[3]), which together imply (c) *equiprobable divisions* (Felsenthal and Machover 1998, 37-38; cf. Penrose 1946, 53). The typical justification for these assumptions is that, because a priori power is power solely in virtue of the formal-institutional rules, measuring it requires abstracting from information related to factors such as the distribution of preferences. Once such information is excluded, equiprobable divisions are supposed to be justified by the principle of insufficient reason.[4]

A principal objection against the concept of a priori voting power is that, by abstracting from the actual distribution of preferences, it is based on decisiveness just as much in logically possible but practically impossible divisions as in actually likely ones and, consequently, overestimates the power of actors with extreme preferences (Garrett and Tsebelis 1999). It ignores, in other words, the voting power of actors based on their position, not in the formal-institutional voting structure, but in other social structures that shape preferences. A posteriori voting power incorporates this further information, and we can generalize *PB* to yield a measure equipped to take into account this information and the consequent probability $p(d_j)$ of each division d_j ($\sum p(d_j)=1$). The *Generalized Penrose-Banzhaf Measure* (*GBP*) replaces [W.1] with:[5]

[W.2] $$w(d_j) = p(d_j)$$

PB thus turns out be a special case of *GPB*: *PB*'s a priori assumption of equiprobable divisions amounts to assuming that $p(d_j) = \frac{1}{|D|}, \forall j$, whereas *GPB* yields an a posteriori measure if $p(d_j)$ is based on information about the actual distribution of, and correlations between, preferences. In voting structures with more than two alternatives, because voters may have incentives to vote strategically for a less preferred alternative, $p(d_j)$ may also reflect the potential effects of strategic interaction on behaviour. The *SS* index, interpreted probabilistically, also turns out to be a special a posteriori case of *GPB* (see Appendix).

A posteriori measures, however, face a potential conceptual challenge stemming from the conditional-dispositional nature of agential power (Morriss 2002). A person does not suddenly become more powerful by restricting her preferences to the few outcomes she is able to effect: she remains powerless to effect the other outcomes, even though she no longer prefers to. Conversely, someone able to effect outcomes she does not actually prefer may be less inclined to effect them, but nevertheless has the power to do so. Thus in assessing an individual's agential power, we are asking a counterfactual question from an ex ante perspective, something like: *Would* she be able to effect these outcomes *if* she wanted to (Goldman 1972, 223)? We are not asking an ex post question about what actually happened, such as: *Was* she able to effect the outcomes she actually preferred? The measurement of voting power requires counterfactualizing over the outcome-preferences of the power-assessee (Braham and Holler 2005, 139).[6]

Therefore, in measuring v_i's voting power, we may base $p(d_j)$ on information about other voters' preferences, and correlations between preferences (including the power-assessee's), but not the content of v_i's preferences. We may therefore calculate voting power by dropping the equiprobable voting assumption for other voters, and dropping the voting independence assumption, but must retain the equiprobable voting assumption for the power-assessee. This models the power-assessee's ex ante first-person perspective in deliberating about what to do, from which her preferences and actions (albeit not the factors shaping them) are an open question.

Key Assumptions of the Classic Model

The (Generalized) Penrose-Banzhaf measure rests on two key microfoundational assumptions. To articulate these assumptions, we should first distinguish between agential power, which may or may not be exercised, and its efficacious exercise. An agent may have the power to effect an outcome but choose not to (because she prefers the alternatives). This is why agential power cannot be equated with actually effecting outcomes: it is a conditional-dispositional concept. We analyze power from an ex ante perspective, but its efficacious exercise from an ex post perspective.

The first assumption is to conceive power in terms of hypothetical efficacy. The second is that an actor's (ex ante) voting power is a function of the power she hypothetically would have efficaciously exercised ex post in each logically possible division were that division to obtain (weighted by the division's probability). This assumption is necessary to justify calculating voting power in a voting structure as a whole ex ante on the basis of individual events (divisions) within the space of possibilities (all possible divisions)—where for each division the voter is assigned a division efficacy score reflecting the extent to which she would have efficaciously exercised power to effect the outcome ex post. These individual divisions—more specifically, division efficacy scores—furnish the microfoundations of the notion of voting power as measured by *GPB*.

Any member of the family of Efficacy Measures requires both (*efficacy* and *ex-ante/ex-post*) assumptions. The classic Penrose-Banzhaf model, however, adds another: the *all-or-nothing* assumption. This reflects a particular conception of efficacy, according to which an actor is efficacious for an outcome if and only if she has at least one other action-strategy such that, if she undertook it instead, the outcome would be different (holding everyone else's actions fixed). This conception of efficacy as (full) decisiveness equates it with being *necessary* for the outcome. It implies a voter is either fully efficacious in effecting the outcome, or not at all—which is reflected in the fact that, on the classic model, [E.1] $e_i(d_j)$ is either 1 or 0.

The Classic Model's Implausible Microfoundations

The all-or-nothing assumption saddles the measure of voting power with implausible microfoundations. It builds the (ex ante) measure of voting power within a voting structure on the basis of a mistaken analysis of the power hypothetically exercised (ex post) by voters to effect outcomes in individual divisions. In particular, *overdetermined* outcomes are erroneously deemed to have not been effected by the individual agential power of any particular voter.

Begin by considering three voters under dictator-rule voting (D3) and under random-rule voting (R3) (where voters vote but outcomes are selected randomly). In D3, v_1 is decisive in all divisions; no other voter is decisive in any division. While the dictator has the maximum voting power possible (PB_1=1), the dummies have zero. In R3, by contrast, every voter is a dummy and Nature a dictator: *no* voter is decisive in any of the eight divisions. If we restrict attention to D3 and R3, the classic model's microfoundations seem plausible. In each division of D3, the outcome is construed as having been (hypothetically) effected ex post by v_1 but not v_2 or v_3, i.e., in each division, only v_1 is construed as efficaciously exercising voting power. This is plausible because it explains the outcome's relation to the voters as a whole and to individual voters separately. These two types of fact are also explained for R3, where outcomes are not effected by any voter's power.

M3 and U3, by contrast, expose the mistake underlying the classic model. Consider division d_1 of M3 in Table 1, where each voter has voted YES. Because no voter is decisive for the outcome, the classic model gives $e_i(d_1)$=0 to each. It assumes that ex post *no* voter has effected the division outcome. But this is false: unlike divisions of R3, each outcome is a function of how the voters have voted. Rather than no voter causally effecting the outcome, this is a case in which the outcome is *overdeterminately* effected by the exercise of their voting power.

The classic model fares no better in U3. Under unanimity-rule voting, each voter has a veto over the YES-outcome. This implies, conversely, that each can unilaterally effect a NO-outcome. Now consider d_8 of U3 in Table 2, where everyone votes NO. Once again, because no voter is decisive, the classic model gives $e_i(d_8)=0$ to each voter: it assumes that ex post the division outcome is not causally effected by any voter. Again, this is false: unlike R3, the outcome here is causally determined by the voters—overdetermined, not non-determined.

Joint Intentional Action and Distributive Power-Blockage

The only hope for salvaging the classic model would be to argue that while overdetermined outcomes are *collectively* effected by some group of voters as a whole, it does not follow that any *particular* voter efficaciously exercises power to effect them (Barry 2002, 181-82). Sometimes a group's collective power may preclude distribution amongst individual members, i.e., sometimes individuals may not themselves have or exercise agential power in virtue of their membership of a group with collective power. Call this phenomenon *distributive power-blockage* (Abizadeh 2021b, 310-312).

This attempt to salvage the classic model requires vindicating two theses: that distributive power-blockage exists; and causally overdetermined divisions always exhibit it. If power-blockage never obtains, then the classic model could not explain how a group could jointly exercise power to effect outcomes even though no particular member does; and if only *some* overdetermined outcomes exhibit power-blockage, then the classic model could not explain those that do not. I argue that, even if power-blockage could be vindicated for some cases, it cannot be done for all.

The premise that power is jointly exercised by a group, but not individually by its members, suggests that the only plausible candidates for power-blockage—for collective power without individual agential power—are cases of joint intentional action. To stave off the objection that any group's collective power distributes amongst the individuals who control its collective decision-making process (e.g. a group agent (List and Pettit 2011) such as a political party controlled by its leaders), consider a group lacking any such process. Imagine we are 100 inhabitants each of whom knows that, without many of us contributing to block the overflowing river, our neighbourhood will be flooded. Each intensively wants to avoid disaster, but we lack any joint decision-making process or even any means to forge an explicit agreement to act together. How could such a group undertake joint intentional action to prevent flooding?

Imagine each of us contributes on the following basis: each believes that, if enough people contribute, the flooding will be prevented; each intends to contribute conditional on many others doing so as well; each expects many others will contribute; and all this is common knowledge. On one view, strategic coordination on the basis of common knowledge such as this can qualify as the joint action of a group (Chant and Ernst 2007). Bratman (2014, 43, 92-95) argues, by contrast, that strategic coordination cannot amount to joint intentional action, because merely intending that one contribute, on the *expectation* that others will, does not yield a joint intention. We can adjust our example to accommodate this more demanding view by assuming that each intends not just that *she* contribute, but that *we* each contribute—which can yield joint intentional action, on Bratman's view, because, in light of rational norms of planning, such an intention implies (as strategic coordination does not) the disposition to support others' contributions (cf. Searle 2010).

If either of these accounts of joint intentional action succeeds, we have a potential candidate for collective power with power-blockage. Now imagine that although all 100 individuals contribute, 80 would have been sufficient. This implies no one's individual contribution was necessary: the other 99 would have prevented the flooding on their own. One might therefore conclude that, while the inhabitants exercised collective power qua group, no one in particular exercised any individual power in effecting the outcome (because, by hypothesis, no one would have been able unilaterally to alter which outcome the group effected).

Now assume arguendo that this power-blockage argument could be vindicated in some such case. The problem is that this would still be insufficient to salvage the classic model, because there also exist overdetermined divisions without joint intentional action. Consider overdetermined divisions in voting structures without strategic interaction. The simple voting games we have been considering qualify: with only two alternatives, each voter sincerely votes her preferred option regardless of what others do. Now all we need do is further stipulate that voters do not have joint intentions in Bratman's stronger sense either (and do not form a group agent). Even in these circumstances—in which there is no question of joint intentional action and hence no group whose collective power could be non-distributed—the classic model assigns $e_i(d_j)$=0 to all voters in overdetermined divisions such as d_1 of M3 and d_8 of U3, falsely implying that voters fail to exercise any power in effecting the outcomes.

The classic model, in other words, runs afoul of overdetermined divisions in which individuals just happen to vote such that their votes together effect an outcome. The power-blockage argument—even if such a phenomenon existed—is unavailable in these cases. The classic model recognizes that in a division such as d_5 of M3, in which v_1 and v_2 efficaciously exercise power with each other to effect a YES-outcome, v_1 and v_2's collective power in effecting the outcome distributes amongst v_1 and v_2. But it fails to recognize that the same holds in cases of overdetermination—with or without joint intentional action. These are cases of *over*determination, not individual *non*-determination.

Erroneous Outcomes of the Classic Model

The classic model does not merely build on mistaken analyses of individual divisions at the microfoundational level. It also yields erroneous measures of voting power, as becomes manifest once we drop the voting independence assumption. Take the binary majority-rule voting structure M5={3; 1,1,1,1,1}, with $|D|=2^5=32$ possible divisions. Consider the case of No Close Calls, where $p(d_j)=0$ for all 20 divisions in which the outcome wins by 3-2, and $p(d_j)=\frac{1}{12}$ for all twelve remaining divisions. Since no voter is ever decisive in any of the twelve positively weighted divisions, the classic model yields $GPB_i=0$ for all voters. In other words, the classic model does not merely presuppose that, ex post, no one has efficaciously exercised power in 5-0 and 4-1 divisions. It also implies that, given the existing preference-shaping social structures in No Close Calls, no voter has any a posteriori voting power to effect outcomes ex ante (cf. Machover 2007, 3-4; Diskin and Koppel 2010, 109). This is erroneous because in this voting structure—unlike random rule—outcomes are determined by how voters vote.

It might be objected that the result in No Close Calls stems not from the all-or-nothing assumption, but from *GPB*'s purported adherence to a widespread Weberian thesis about social power, namely, that social power presupposes the capacity to overcome resistance (Weber 1978, 53; Goldman 1974, 231; Barry 1988; Braham 2008). On this basis, one might conclude that although voters in d_1 of M3, in which all three voters vote YES, do indeed (overdeterminately) *cause* the outcome, they do not efficaciously exercise social power, because they do not face any resistance, and that *GPB* therefore correctly counts this as a division in which no voter efficaciously exercises (social) power.

The conclusion rests on two mistakes. First, the resistance thesis does not state that having (or efficaciously exercising) social power presupposes *actual* resistance (or overcoming it). It states, rather, that social power presupposes either actual or *hypothetical* resistance: that one could overcome were others to resist (Braham 2008, 15). So the failure to assign any efficacious power to voters in d_1 of M3 cannot be justified by the resistance thesis: although v_1 helps cause the YES-outcome without actually overcoming anyone's resistance, she could (thanks to v_3's assistance) help overcome v_2's resistance were v_2 to vote NO (as indicated by d_3).

Second, *GPB* does not in any case adhere to the resistance thesis. Consider d_1 of U3 in which all three voters vote YES. The classic model sets $e_1(v_i)=1$ for each voter, correctly assuming each voter efficaciously exercises power to effect the division outcome ex post. Yet not only is there no actual resistance, in U3 no one could even in principle overcome *anyone*'s resistance to a YES-outcome: if anyone were to resist, the outcome would be NO. The only voting power voters here have to effect a YES-outcome is a power-with-others that is not power-despite anyone.[7]

Machover (2007, 4) suggests *GPB* misfires in No Close Calls perhaps because it presupposes a concept of power in some sense too individualistic. This is on the right track, but it would be mistaken to deem the classic model individualistic in the sense that it implicitly reduces social power to unilateral power—to the exclusion of power-with. True, in cases of overdetermination—including the twelve positively weighted divisions of M5 under No Close Calls—voters effect the outcome *with* each other, and the classic model sets $e_i(d_j)=0$ there. But the model does count many instances of acting with others' assistance as occasions in which a voter has efficaciously exercised power. Consider d_5 of M3, in which v_1 and v_2 vote YES and v_3 votes NO. Both v_1 and v_2 are decisive in this division, and *GPB* assigns each of them $e_1(d_5)=e_2(d_5)=1$. Yet v_1 could not have efficaciously exercised power in this division without v_2's assistance (if v_2 had voted NO, v_1 would not have been successful). Although *GPB* does not distinguish between unilateral power and power-with, it includes both. The correct diagnosis lies elsewhere.

What generates the mistaken verdict in No Close Calls is the classic assumption that efficacious power stems only from *full efficacy*. What is needed to explain the overdetermined divisions (at the microfoundational level), and to yield plausible results in cases such as No Close Calls (at the output level), is a conception of efficacy that admits of degrees. The all-or-nothing assumption does indeed reflect an individualistic bias: not that it forces *GPB* to ignore power-with, but that it imposes an unduly individualistic conception of an *individual's* power-with in virtue of (potentially) acting with a group with collective power. The classic model assumes an individual efficaciously exercises power-with-others only if she could unilaterally scuttle the group's collective power. But that is false, and we can see why by recognizing *partial* efficacy.

Degrees of Efficacy

Before explicating the concept of partial efficacy, two preliminary points. First, I use partial efficacy and partial decisiveness as equivalent labels for the same concept. For some, decisiveness connotes an all-or-nothing property, so to avoid controversy, one could stick with efficacy for the more general concept. But nothing substantive turns on the labelling question.

Second, I prepare the way for examining the concept by showing how it could be represented. Consider three voters using lottery-rule voting (L3), in which outcomes are determined by a single, randomly selected voter's vote (Amar 1984). Let r_i be the probability of v_i's vote being randomly selected to determine the outcome ($\sum r_i$=1), which is also the probability v_i's vote will be decisive in a given division. Equiprobability lottery-rule voting EL3 is a special case of L3 in which $r_i = \frac{1}{|V|}, \forall i$. To represent EL3, we could, as Table 3 does, represent Nature as a fictional voter v_0 with three possible moves

Table 3: Lottery-rule Voting EL3

Divisions$^+$ D^+ (|D|$^+$=3·2^3=24)	$a_0(d_j)$	$a_i(d_j)$, $e_i(d^+_{j,k})$			Out-come
		v_1	v_2	v_3	
$d^+_{1,1}$	$a_1(d_1)$	y,1	y,0	y,0	y
$d^+_{1,2}$	$a_2(d_1)$	y,0	y,1	y,0	y
$d^+_{1,3}$	$a_3(d_1)$	y,0	y,0	y,1	y
$d^+_{2,1}$	$a_1(d_2)$	**n**,1	y,0	y,0	n
$d^+_{2,2}$	$a_2(d_2)$	n,0	y,1	y,0	y
$d^+_{2,3}$	$a_3(d_2)$	n,0	y,0	y,1	y
$d^+_{3,1}$	$a_1(d_3)$	y,1	n,0	y,0	y
$d^+_{3,2}$	$a_2(d_3)$	y,0	**n**,1	y,0	n
$d^+_{3,3}$	$a_3(d_3)$	y,0	n,0	y,1	y
$d^+_{4,1}$	$a_1(d_4)$	**n**,1	n,0	y,0	n
$d^+_{4,2}$	$a_2(d_4)$	n,0	**n**,1	y,0	n
$d^+_{4,3}$	$a_3(d_4)$	n,0	n,0	y,1	y
$d^+_{5,1}$	$a_1(d_5)$	y,1	y,0	n,0	y
$d^+_{5,2}$	$a_2(d_5)$	y,0	y,1	n,0	y
$d^+_{5,3}$	$a_3(d_5)$	y,0	y,0	**n**,1	n
$d^+_{6,1}$	$a_1(d_6)$	**n**,1	y,0	n,0	n
$d^+_{6,2}$	$a_2(d_6)$	n,0	y,1	n,0	y
$d^+_{6,3}$	$a_3(d_6)$	n,0	y,0	**n**,1	n
$d^+_{7,1}$	$a_1(d_7)$	y,1	n,0	n,0	y
$d^+_{7,2}$	$a_2(d_7)$	y,0	**n**,1	n,0	n
$d^+_{7,3}$	$a_3(d_7)$	y,0	n,0	**n**,1	n
$d^+_{8,1}$	$a_1(d_8)$	**n**,1	n,0	n,0	n
$d^+_{8,2}$	$a_2(d_8)$	n,0	**n**,1	n,0	n
$d^+_{8,3}$	$a_3(d_8)$	n,0	n,0	**n**,1	n
$PB_i = \frac{\sum_{d^+_{j,k}\in D^+} e_i(d^+_{j,k})}{\|D^+\|}$		⅓	⅓	⅓	*PB*(*G*) = 1

(picking the vote of one actual voter). Calling the combination of division d_j plus Nature's move a *division*$^+$ (represented as $d^+_{j,k}$), we then calculate PB_i on the basis of v_i's *division*$^+$ *efficacy scores* $e_i(d^+_{j,k})$. Alternatively, we could, as in Table 4, represent Nature's move by assigning fractional values to voters' division efficacy scores. That is, rather than calculating PB_i on the basis of v_i's division$^+$ efficacy scores, where $e_i(d^+_{j,k})$=1 if v_i is decisive in $d^+_{j,k}$ and 0 if not, we calculate PB_i as before on the basis of *division* efficacy scores, but set $e_i(d_j)=r_i$ for each division.

The significance of Table 4's mode of representation is that, in assigning fractional values to $e_i(d_j)$, it technically abandons the all-or-nothing assumption. However, the abandonment is purely technical: it does not invoke a conception of efficacy that admits of degrees. The fraction represents the ex ante *probability* of being (fully) decisive in a division, not the voter's partial decisiveness in the ex post sense of only partially having efficaciously exercised power. Ex post, one has either won the lottery and fully determined the outcome, or not. To see how fractional values could represent partial efficacy, we must turn to causally overdetermined outcomes.

Table 4: Lottery-rule Voting EL3, with fractional division efficacy scores

Divisions D ($\|D\|=2^3=8$)	$a_i(d_j)$, $e_i(d_j)$			Outcome
	v_1	v_2	v_3	
d_1	y,⅓	y,⅓	y,⅓	p(y)=1, p(n)=0
d_2	n,⅓	y,⅓	y,⅓	p(y)=⅔, p(n)=⅓
d_3	y,⅓	n,⅓	y,⅓	p(y)=⅔, p(n)=⅓
d_4	n,⅓	n,⅓	y,⅓	p(y)=⅓, p(n)=⅔
d_5	y,⅓	y,⅓	n,⅓	p(y)=⅔, p(n)=⅓
d_6	n,⅓	y,⅓	n,⅓	p(y)=⅓, p(n)=⅔
d_7	y,⅓	n,⅓	n,⅓	p(y)=⅓, p(n)=⅔
d_8	n,⅓	n,⅓	n,⅓	p(y)=0, p(n)=1
PB_i	⅓	⅓	⅓	$PB(G)$ = 1

Begin by considering three non-voting cases without overdetermination:

Unique Assassin: The victim has taken refuge in a location known and accessible to no one else, except one unique assassin. If not for the unique assassin, the victim would live until old age. But the assassin accesses the location and shoots the victim, killing him.

Lone Assassin: Many assassins could kill the victim. But a lone assassin shoots and kills the victim (because only he has a motive to do so).

Two Small Assassins: The victim is tied up in the trunk. Neither of two assassins is strong enough to push the car over the cliff, but together they can. Both push, killing the victim.

In these cases, each assassin's causal efficacy ex post is explicable along the lines of the classic Penrose-Banzhaf model. The unique assassin's action is both *sufficient* and *robustly necessary* (i.e., necessary on every occasion) for the killing: he is the analogue of the dictator in any division of D3. The lone assassin is analogous to v_1 in d_2 of U3, in which v_1 unilaterally effects a NO-outcome by exercising her veto: his action is *sufficient* and *necessary* for the killing. Finally, the two small assassins are analogous to YES-voters v_1 and v_2 in d_5 of M3: while not sufficient, each assassin's action is *necessary*. These cases are all adequately dealt with by the classic model, because decisiveness is equivalent to necessity, which is present in each case.

Now consider cases of overdetermination (McDermott 1995; Ramachandran 1997; Hitchcock 2001; Schaffer 2003; Braham 2008) analogous to the divisions whose overdetermined outcomes the classic model assumes are *not* effected by the power exercised by any voter.

Three Small Assassins: The victim is tied up in the trunk. None of the three assassins is strong enough to push the car over the cliff, but two are sufficient. All three push, killing the victim.

Because this is a case of causal overdetermination, no single assassin was *necessary* for the victim's death: had any one not pushed, the pushing by the other two would have been sufficient. Nor was any assassin's pushing *sufficient*: had only one pushed, the victim would have lived. The three small assassins are like the three voters in d_1 of M3, each of whom has voted YES, and whom the classic model portrays as not having exercised any power to effect the outcome. Yet the killing was obviously caused by the three assassins.

Two Eager Assassins: Two assassins independently and simultaneously shoot, killing the victim. Each assassin's shot was sufficient for the kill.

Here, each assassin's shot was *sufficient*, yet, because the outcome is overdetermined, neither was necessary: had a given assassin not shot, the victim would have still been killed. The two eager assassins are like voters in division d_8 of U3 in Table 2, each of whom votes NO, yet none of whom is necessary for the NO-outcome—and whom the classic model consequently portrays as having not efficaciously exercised power. Yet clearly the victim was killed by the two assassins, just as the NO-outcome was effected by the voters in d_8 of U3: the two assassins efficaciously exercised their power to kill the victim.

We have already ruled out the retort that the group exercises collective power that fails to distribute amongst its members. (Recall: sometimes there is no group to which collective power could be attributed). What of an analogous retort centred on causation, namely, that the mereological sum of the individuals collectively (but no individual separately) causes the outcome (Lewis 1986, 212)? As Schaffer (2003) has argued, the supposition that a collective has emergent causal properties even though no individual member plays a causal role cannot be sustained because (amongst other things) it cannot explain why the collective would lose those causal properties if stripped of its (supposedly causally inert) members. What is needed is an explanation of members' causal role in spite of the fact the outcome does not counterfactually turn on any individual's causal contribution in particular.

How can we explain the causal efficacy of—and power efficaciously exercised by—each assassin in Three Small or Two Eager Assassins? Sufficiency cannot explain the former case; necessity cannot explain either case. (And joint intentional action is clearly missing in the latter.) One way to explain each actor's causal contribution, and hence the causal power each efficaciously exercises, is that each actor's action-event not only is an element of the set of actions-events that together effects the outcome, but also individually satisfies a NESS test (Wright 1985, 1988): each action-event is a *n*ecessary *e*lement of a *s*ufficient *s*et of conditions for the outcome to obtain, which set is a subset of the actually obtaining set of conditions sufficient for the outcome.[8]

To explain: in each case, take the actual action-events and outcome. (An *action-event* is the action indexed to the actor, i.e., the event $ae_i(d_j)=(v_i, a_i(d_j))$ of v_i undertaking action $a_i(d_j)$ in a division or complete action configuration d_j. In what follows, expressions such as v_i's action refer to the action-event.) *Ex hypothesi*, the actual set of action-events is sufficient for the outcome under the circumstances; hence the group's actions compose a *sufficient set of conditions* or SSC for the outcome. Assume each action-strategy can be ranked from least to most favourable to a given outcome, call any actor who resists the actual outcome to the highest degree she can, a *maximal resister*, and assume that maximally resistant action-events make no contribution to effecting the outcome. Call any subset of the action-events of an SSC excluding maximally resistant ones, and that is still sufficient to effect the outcome, a *sufficient set of action-events* for the outcome (sufficient given the non-action circumstances), or SSA. Call the SSA of the SSC that actually obtains in a given action configuration (division) and which includes all of the actual SSC's action-events other than maximally resistant ones, the *actual* SSA.

Now take any subset of the actual SSA that is itself also an SSA (sufficient to effect the outcome, independently of others' actions), but that is *vulnerable* to at least one of its action-elements, in the sense that without it, the rest of the action-events composing the hypothetical SSA would be insufficient to effect the outcome. Call action-events to which an SSA is vulnerable *critical* elements. Without the critical actions-events, a vulnerable SSA, or VSSA, would be insufficient to effect the outcome: critical action-events are necessary for the VSSA to be sufficient for the outcome. Any actor whose action is an element of the actual SSA that effects the outcome, and whose action is also a critical element of a subset of the actual SSA, satisfies the NESS test.

To illustrate: in the Three Small Assassins, let $d_1=\{ae_1(d_1), ae_2(d_1), ae_3(d_1)\}$ be the complete action configuration in which each assassin pushes the car. No assassin's action is critical to the actual SSA. But the actual SSA contains three proper subsets that are SSAs ({$ae_1(d_1)$, $ae_2(d_1)$}, {$ae_1(d_1)$, $ae_3(d_1)$}, {$ae_2(d_1)$, $ae_3(d_1)$}), and each assassin's action is critical in two of them. Similarly, consider Two Eager Assassins: neither assassin's action is critical to the actual SSA, but the actual SSA contains two proper subsets that are SSAs, and each assassin's action is critical in one of them. In both cases each assassin's action therefore satisfies the NESS test and thus, assuming there are no pre-empting causes, is (partially) causally efficacious.

To determine which actors are partially efficacious, one could resort to a shortcut test. Let a *minimal* SSA, or MSSA, be a VSSA vulnerable to all of its action-events. Since an action-event is a critical element of a subset of the actual SSA if and only if it is an element of an MSSA subset of the actual SSA, the shortcut NESS test simply checks whether an individual's action is an element of any such MSSA. If it is, then the actor is at least partially efficacious; if not, then not, and there is no need to check remaining subsets.

But the regular NESS test does not merely establish *which* individuals' action helps cause an outcome; it also serves to determine the *degree* to which it does. To cut to the chase: although in Three Small Assassins v_1's action is not a necessary and hence fully efficacious cause, it is *partially efficacious*; in particular, in virtue of being a critical element of two of the three hypothetical SSA proper subsets of the actual SSA, v_1's action is ⅔ efficacious. Similarly, although in Two Eager Assassins v_1's action is not necessary, in virtue of being a critical element of one of the two hypothetical SSA proper subsets of the actual SSA, v_1's action is ½ efficacious. In some cases (such as these two), the shortcut yields the same degree of efficacy but, as we shall see, in others it does not.[9]

Voting Power with Partial Efficacy

I propose a similar analysis of efficacy for ascertaining voting power. Just as what is needed to explain each assassin's causal efficacy in casually overdetermined deaths is a concept of *degrees* of causal efficacy, so too what is needed to explain voters' efficacious exercise of power in overdetermined divisions is the notion of partial efficacy.

Tuck (2008) has suggested something along these lines. Tuck denies that no individual efficaciously exercises causal power in overdetermined divisions lacking decisive voters. His argument relies on the fact that in any division, of the group of actors whose actions compose the actual SSA, there is always at least one subset of actors whose actions compose an MSSA vulnerable to *each* of its action-events. Tuck assumes that of the MSSA subsets of the actual SSA, one and only one—the actions of voters whom he calls the efficacious set—has causal efficacy over the outcome, and that the action of *each* member of this efficacious set has full causal weight, while the actions of everyone else have none. He then argues that since, in any division lacking decisive voters, for any voter who voted in favour of the outcome, we cannot know whether she is in the efficacious set, the question is what the probability is she is. His answer is that for any successful voter, the probability her action has fully caused the outcome is equal to the proportion of all MSSAs, which are a proper subset of the actual SSA, of which the voter's action is an element (2008, 40-45). Take division d_1 of M3, in which voters unanimously vote YES. The actual SSA consists in $\{ae_1(d_1), ae_2(d_1), ae_3(d_1)\}$, which has three proper MSSA subsets: $\{ae_1(d_1), ae_2(d_1)\}$, $\{ae_1(d_1), ae_3(d_1)\}$, $\{ae_2(d_1), ae_3(d_1)\}$. Tuck supposes that one of these three MSSAs causes the outcome; he would therefore conclude that the probability v_1 was fully causally efficacious in producing the outcome is ⅔, i.e., $e_i(d_1)=⅔$ for each voter in d_1 of M3.

Talk of probability suggests Tuck is implicitly deploying a principle of insufficient reason as warrant, under conditions of total ignorance, for attributing to each logically possible MSSA proper subset of the actual SSA an equal probability of containing the actions of the efficacious set. However, the meaning of probability is unclear in this context. Unlike lottery voting, a fractional $e_i(d_j)$ here does not represent the probability of being fully efficacious—no one is decisive. It represents, rather, the fact that a successful voter whose action satisfies a NESS test in an overdetermined division is partially efficacious. The most plausible way to interpret Tuck's approach is to translate his talk of probability of full efficacy into partial efficacy.

With this adjustment, we can formalize a Tuckian approach to partial efficacy. In a given voting structure, let a division be the *child* of another division (the *parent*) if and only if both are identical except that the vote of one voter in the former is less favourable to the outcome of the latter division's outcome. For each division d_j, let LC_j be its set of *loyal children*, comprising those children whose outcome is the same as their parent division; LD_j its set of *loyal descendants*, comprising its loyal children, their loyal children, etc.; and MLD_j its set of *minimal* loyal descendants whose actual SSA is an MSSA. A voter's partial division efficacy score on the Tuckian approach is equal to the proportion of minimal loyal descendants in which she is decisive:

[E.2] $$e_i(d_j) = \begin{cases} 1 & \text{if } v_i \text{ is decisive in division } d_j \\ 0 & \text{if } v_i \text{ is not decisive in division } d_j \text{ and } d_j \text{ has no loyal descendants} \\ \frac{1}{|MLD_j|}\sum_{d_k \in MLD_j} e'_i(d_k) & \text{otherwise} \end{cases}$$

where

[E.2.1] $$e'_i(d_k) = \begin{cases} 1 & \text{if } v_i \text{ is decisive in division } d_k \\ 0 & \text{otherwise} \end{cases}$$

Take the division of M5 with four YES-voters: $e_i(d_j)$=¾ for each YES-voter, because this division has four minimal loyal descendants and each YES-voter is decisive in three of them.

The Tuckian approach has two flaws. First, it implicitly uses the shortcut NESS test, to determine not merely which voters are (partially) efficacious, but also the degree to which each is. The shortcut yields the same $e_i(d_j)$ as the regular NESS test in the kinds of divisions Tuck considers—cases in which every vulnerable loyal descendant is minimal. But it yields mistaken measures in cases containing non-minimal vulnerable loyal descendants.

Consider the following binary voting structure (Braham and van Hees 2009), with five voters deploying the following voting rule: a YES-outcome results if and only if either v_1 votes YES or at least three other voters vote YES. This can be represented by the weighted voting structure {3;3,1,1,1,1}, which can be thought of as modelling a bicameral legislature in which v_1 composes one chamber and v_2–v_5 compose another, and in which a majority in either chamber is necessary and sufficient to pass legislation.) This voting structure has 32 possible divisions. Consider d_1, in which all five voters vote YES and which has 20 loyal descendants (Table 7). We use the symbol ^{v} to indicate loyal descendants whose actual SSA is a VSSA, and ^{m} for minimal loyal descendants (whose actual SSA is an MSSA).

Table 7: Loyal Descendants of Unanimous YES-Division in {3;3,1,1,1,1}

Loyal Descendants of d_1 (|LD$_1$|=20)	$a'_{1,j}(v_i)$, $e'_{1,j}(v_i)$					Outcome
	v_1	v_2	v_3	v_4	v_5	
d_2	n,0	y,0	y,0	y,0	y,0	y
d_3	y,0	n,0	y,0	y,0	y,0	y
d_4^{m}	n,0	n,0	$\underline{y}$,1	$\underline{y}$,1	$\underline{y}$,1	y
d_5	y,0	y,0	n,0	y,0	y,0	y
d_6^{m}	n,0	$\underline{y}$,1	n,0	$\underline{y}$,1	$\underline{y}$,1	y
d_7^{v}	$\underline{y}$,1	n,0	n,0	y,0	y,0	y
d_9	y,0	y,0	y,0	n,0	y,0	y
d_{10}^{m}	n,0	$\underline{y}$,1	$\underline{y}$,1	n,0	$\underline{y}$,1	y
d_{11}^{v}	$\underline{y}$,1	n,0	y,0	n,0	y,0	y
d_{13}^{v}	$\underline{y}$,1	y,0	n,0	n,0	y,0	y
d_{15}^{v}	$\underline{y}$,1	n,0	n,0	n,0	y,0	y
d_{17}	y,0	y,0	y,0	y,0	n,0	y
d_{18}^{m}	n,0	$\underline{y}$,1	$\underline{y}$,1	$\underline{y}$,1	n,0	y
d_{19}^{v}	$\underline{y}$,1	n,0	y,0	y,0	n,0	y
d_{21}^{v}	$\underline{y}$,1	y,0	n,0	y,0	n,0	y
d_{23}^{v}	$\underline{y}$,1	n,0	n,0	y,0	n,0	y
d_{25}^{v}	$\underline{y}$,1	y,0	y,0	n,0	n,0	y
d_{27}^{v}	$\underline{y}$,1	n,0	y,0	n,0	n,0	y
d_{29}^{v}	$\underline{y}$,1	y,0	n,0	n,0	n,0	y
d_{31}^{m}	$\underline{y}$,1	n,0	n,0	n,0	n,0	y

Because only five loyal descendants of d_1 contain an actual SSA that is an MSSA, and v_1's action is an element of only one of these (d_{31}), and each of the other voters' actions is an element of three, on the Tuckian approach $e_1(d_1)$=⅕ and $e_{2\text{–}5}(d_1)$=⅗. Thus the voter with the *greatest* weight is counted as the *least* efficacious. The problem stems from using the shortcut NESS test. The Tuckian approach fails to take into account loyal descendants in which v_1 is decisive when other voters (impotently) vote in favour of the outcome that v_1 effects unilaterally (such as d_7).

To correct for this mistake, we must return to the regular NESS test. The following equates $e_i(d_j)$ with the proportion of loyal descendants (not just minimal ones) for which v_i is decisive:

[E.3] $$e_i(d_j) = \begin{cases} 1 & \text{if } v_i \text{ is decisive in division } d_j \\ 0 & \text{if } v_i \text{ is not decisive in division } d_j \text{ and } d_j \text{ has no loyal descendants} \\ \frac{1}{|LD_j|}\sum_{d_k \in LD_j} e'_i(d_k) & \text{otherwise} \end{cases}$$

where

[E.3.1] $$e'_i(d_k) = \begin{cases} 1 & \text{if } v_i \text{ is decisive in division } d_k \\ 0 & \text{otherwise} \end{cases}$$

This approach, roughly along the lines of Braham and van Hees's approach to measuring degrees of causal efficacy,[10] yields more plausible results for d_1 of {3; 3,1,1,1,1}. Because v_1 is decisive in 11 of d_1's 20 loyal descendants, and each other voter in three, this new approach yields $e_1(v_1)=\frac{11}{20}$ and $e_1(v_{2\text{–}5})=\frac{3}{20}$, which correctly assigns relatively greater efficacy to v_1.

The Tuckian approach's second flaw, however, also afflicts the Braham-van-Heesian approach: both drop the classic all-or-nothing assumption for divisions (in [E.2] and [E.3]), but retain it for their loyal descendants, setting descendants' efficacy scores to either 0 or 1 (in [E.2.1] and [E.3.1]). They therefore count overdetermined loyal descendants in which the actual SSA is not a VSSA as not caused by any voter. This is incoherent: the justification for degrees of efficacy in divisions applies equally to their descendants. The value of $e'_i(d_k)$ for a division d_k considered as a descendant should match the value $e_i(d_j)$ of the identical division (j=k) considered as a parent.

What we need is a *Recursive Measure* of voting power *RM*, defined as [EM], where, like *GPB*, the division weight is [W.2] $w(d_j) = p(d_j)$, but where a voter's division efficacy score $e_i(d_j)$ is defined recursively, calculated[11] on the basis of not just decisiveness in the division itself, nor merely in its loyal children, but also their loyal children, and so on:[12]

[E.4] $$e_i(d_j) = \begin{cases} 1 & \text{if } v_i \text{ is decisive in division } d_j \\ 0 & \text{if } v_i \text{ is not decisive in division } d_j \text{ and } d_j \text{ has no loyal children} \\ \frac{1}{|LC_j|}\sum_{d_k \in LC_j} e_i(d_k) & \text{otherwise} \end{cases}$$

We can illustrate *RM* by contrasting it with the classic measure *GPB*, which defines $e_i(d_j)$ as in [E.1], and the non-recursive Braham-van-Heesian measure *BH*, which defines $e_i(d_j)$ as in

[E.3]. Consider again M5 assuming equiprobable divisions. Of the 32 divisions, the outcome is effected by (i) 5-0 votes in two divisions $d_{1=}$(y,y,y,y,y) and $d_{32=}$(n,n,n,n,n),[13] (ii) 4-1 in ten divisions, and (iii) 3-2 in 20. In all divisions, $e_i(d_j)$=0 for unsuccessful voters. (i) In the two 5-0 divisions, each voter is successful but none decisive. The classic model [E.1] therefore sets $e_i(d_1)$=$e_i(d_{32})$=0 for each voter. The non-recursive model [E.3], by contrast, assigns fractional values. Each 5-0 division has 15 loyal descendants: $\binom{5}{4}$=5 loyal children won by 4-1 (e.g. (y,y,y,y,n)) and $\binom{5}{3}$=10 loyal grandchildren won by 3-2 (e.g. (y,y,y,n,n)). In each division's ten 3-2 loyal grandchildren, each voter would be decisive in $\binom{4}{2}$=6 of them; in the five 4-1 loyal children, none would be. Therefore, on the non-recursive model [E.3], $e'_i(d_k)$=1 for six loyal descendants, $e'_i(d_k)$=0 for the rest, yielding $e_i(d_1)=e_i(d_{32})=\frac{[6(1)+4(0)]+5(0)}{15}=⅖$. By contrast, on the recursive model [E.4], we first turn to the five 4-1 loyal children, each of which, in turn, begets $\binom{4}{3}$=4 loyal grandchildren won by 3-2, none of whom beget any further loyal descendants.[14] Each voter will be successful in four of the five loyal children, and a voter who is successful in a loyal child will be decisive in three of the four loyal grandchildren it begets, which yields $e_i(d_1)=e_i(d_{32})=\frac{1(0)+4(\frac{1(0)+3(1)}{4})}{5}=⅗$ for the two unanimous divisions. (ii) Now, turning to the ten 4-1 divisions: since no voter is decisive, on the classic model [E.1] $e_i(d_j)$=0 for each; but since each voter who would be successful in a 4-1 division would be decisive in three of its $\binom{4}{3}$=4 loyal children, the non-recursive [E.3] and recursive [E.4] models yield $e_i(d_j)$=¾ for such voters. (iii) Finally, in the 20 3-2 divisions, since each successful voter is decisive, $e_i(d_j)$=1 for each on all three models. The result is that voting power is different according to the three measures. Since every voter is successful in (i) both 5-0 divisions, (ii) eight of the ten 4-1 divisions, and (iii) 12 of the 20 3-2 divisions, $GPB_i=\frac{2(0)+8(0)+12(1)}{32}=0.375$, $BH_i=\frac{2(⅖)+8(¾)+12(1)}{32}=0.5875$, and $RM_i=\frac{2(⅗)+8(¾)+12(1)}{32}=0.6$.

Interpretation, Reasonability, and Puzzle-solving Power of RM

The classic model's strength is that it affords *GPB* an intuitive interpretation: a voter's ex ante *probability* of being decisive. Its shortcoming is that power cannot be analyzed solely in terms of (full) decisiveness. Does *RM* invite an equally intuitive interpretation? Yes. Formally, *RM* represents an *expected value*: the voter's *expected efficacy* in a voting structure. An expected value sums over the value associated with each possible occasion (i.e., a voter's division efficacy score) weighted by the occasion's probability. And intuitively, *RM* captures the extent to which a voter could expect to contribute to effecting the outcome—whether she is decisive or not.

RM not only furnishes the measure of voting power more plausible microfoundations, it also satisfies a set of reasonable postulates for measures of voting power (Abizadeh and Vetta working paper-b). It also solves a number of output-level puzzles faced by *GPB*. Recall the M5 case of No Close Calls, in which $p(d_j)=0$ for 3-2 divisions and $p(d_j)=\frac{1}{12}$ for the rest. The classic model absurdly yields $GPB_i=0$ for all voters—as if voters have no power to effect outcomes in this voting situation. The recursive model solves this puzzle, yielding $RM_i=⅗$.

Both No Close Calls and the equiprobable M5 example above show that the recursive model yields different absolute measures of voting power than the classic model. It can also yield different *ratios* of power between voters—even a priori. For example, in the simple binary three-voter structure in which v_1 has a veto (i.e., proposals are passed if and only if v_1 and at least one of v_2 or v_3 vote in favour), with $p(d_j) = \frac{1}{|D|}, \forall j$, $GPB_1=¾$ and $GPB_{2\text{-}3}=¼$, which implies v_1 has three times greater voting power than either of the others, whereas on the recursive model $RM_1=\frac{17}{24}$ and $RM_{2\text{-}3}=\frac{17}{48}$, which implies v_1's voting power is only 2 times greater.

The recursive model can also yield different *ordinal* rankings of a posteriori voting power. Consider the politically significant case of a persistent minority confronting a persistent majority.

For simplicity, consider majority-rule binary voting with a persistent majority of six and persistent minority of three voters, where each group always vote together but against the other. Since every division is won by 6-3, no one is ever decisive. Hence, on the classic model every voter has equal a posteriori voting power: GPB_i=0 for everyone. On the recursive model, voting power remains zero for persistent-minority members, but because persistent-majority members are always partially efficacious, RM_i>0 for them. (Since there are six persistent-majority members, but only five are needed to secure the outcome, each is decisive in $\frac{5}{\binom{6}{5}}=\frac{5}{6}$ of the six loyal children of each 6-3 division, yielding $e_i(d_j)=\frac{5}{6}$ for every division. Therefore, $RM_i=\frac{5}{6}$.)

For a more dramatic example, which *reverses* ordinal rankings, consider a variant in which the six-voter persistent majority has two members that, on rare occasions, vote with the minority—say, 1% of the time. On the classic model, persistent-minority members will, absurdly, now have *greater* voting power than the four solid persistent-majority members: GPB_i=0 for the latter, but GPB_i=0.01 for the former (because decisive 1% of the time). On the recursive model, by contrast, RM_i rises to $\frac{5}{6}\cdot\frac{99}{100}$=0.825 for the latter, which correctly assigns greater voting power to solid persistent-majority members (and even greater for the two occasional strayers: RM_i=0.835).

Conclusion: Significance of the Recursive Model

The Recursive Measure of voting power is significant for several reasons. First, it alerts us to an aspect of social power overlooked in the classic approach: the power one exercises with others to effect outcomes even when one could not unilaterally scuttle the outcome (Abizadeh 2021b). This can help explain the intuitive sense of empowerment and effective agency—captured by *RM* in the voting context, but of broader significance—that people often have in uniting to act together in social movements. Second, by taking this aspect of power-with into account, *RM* offers a potential solution to the classic voting paradox (Downs 1957), which puzzles over why instrumentally rational voters would bother voting in mass elections when their chances of being decisive is so miniscule: without appealing to moral motivations, or voting's expressive value, the measure allows us to understand the sense in which voters successfully exercise agential power to effect outcomes even when not decisive. Third, the fact that *RM* yields different ratios and even ordinal rankings of voting power than the classical model is of tremendous significance for both normative and empirical analysis—for example, in debates of institutional design about the weight to assign to subunit votes in federal, international, or supranational institutions, or debates about how much decision-making power each exercises under current institutions. Fourth, *RM* not only measures voting power on the basis of plausible microfoundations, it also solves seemingly intractable problems facing the classic model—such as No Close Calls.

Finally, some of these cases, such as the case of persistent minorities, are of tremendous empirical and normative significance. Consider a racially divided polity in which the majority whites invariably vote together against the minority Blacks. Many have the intuition that persistent-minority members, who would be perpetually outvoted by persistent-majority members under majoritarian decision-making, do not merely enjoy less success, but have less substantive decision-making power, despite enjoying formally equal a priori voting power (Jones 1983; cf. Saunders 2010). As we have seen, however, the classic model cannot register this inequality: $GPB_i=0$ for each voter, from whichever group, because under the circumstances no individual is ever decisive. *RM*, by contrast, does register the inequality. Whereas $RM_i=0$ for persistent-minority members (they are never successful in their votes), $RM_i>0$ for persistent-majority members (each is always partially efficacious). Similarly, in a polity divided between three equal-sized camps, two of whom always vote against each other while the third swings between the two, only the recursive model can register the sense in which members of the swing camp enjoy greater power under majoritarian decision-making than the other two. These results have significant normative implications: they suggest, for example, that substantive political equality in decision-making may sometimes require formal-procedural *inequalities* (Abizadeh 2021a).

Appendix: The Shapley-Shubik Index

When Shapley and Shubik (1954) initially introduced their index, they characterized it in terms of being pivotal rather than decisive. A pivotal voter is one who, in an ordered sequence of voters who sequentially vote in favour of an alternative, is the first whose vote secures it regardless of how subsequent voters vote; *SS* was characterized as the probability of being pivotal if all permutations of voters are equiprobable. As subsequent analysis has shown, however, being pivotal is analytically reducible to being decisive (Turnovec et al. 2008). For example, following Felsenthal and Machover (1998, 196-97), SS_i can be defined for simple voting games by defining the division efficacy score $e_i(d_j)$ as in [E.1], and the division weight $w(d_j)$ as follows (where $K_j \subseteq V$ is the set of voters whose votes agree with the outcome in d_j, $v=|V|$, and $k_j=|K_j|$):

[W.3] $$w(d_j) = \frac{(k_j-1)!(v-k_j)!}{2v!}$$

Now, let z be the number of action-strategies available to each player (z=2 in simple voting games, YES or NO) and, for any given action-strategy, let $S_j \subseteq V$ be the subset of all voters in d_j who undertake it, i.e., a *junction* of voters who vote the same. Then *SS* reduces to a special case of *GPB*, yielding a voter's probability of being decisive, if the probability distribution $p(d_j)$ meets two conditions: first, all v+1 logically possible sizes s_j of S_j are equiprobable; second, all junctions of the same size are equiprobable (Laruelle and Valenciano 2005a, 188). This implies the probability of a junction is inversely related to the number of possible junctions of its size, i.e., it is lower the closer its size is to the mean junction size $\frac{v}{z}$ (Weber 1988, 103; Napel 2019, 110). It follows that *SS* yields v_i's probability of being decisive if:

$$p(d_j) = \frac{1}{(v+1)} \cdot \frac{1}{\binom{v}{s_j} \cdot (z-1)^{(v-s_j)}}, \forall j$$

Here, $\binom{v}{s_j}$ is the number of different possible junctions of size s_j; $v–s_j$ is the number of voters outside the junction, for whom z–1 strategies remaining; and $\binom{v}{s_j} \cdot (z-1)^{(v-s_j)}$ is the number of divisions whose junction S_j is the same size. (See Laruelle and Valenciano (2005a, 188), who give the formula for the special case z=2.) This probability distribution would result if voting behaviour were homogeneous in the sense that the probability any voter votes for some arbitrary alternative is the same for all voters and selected from a uniform distribution on [0,1]. *SS* can therefore be interpreted as the probability a voter will be decisive when voters share a common standard by which they judge the desirability of alternatives (Straffin 1977; Leech 1990; cf. Owen 1975).

Indeed, *SS* not only can but arguably must be construed as a special case of a posteriori power *GPB* if it is to interpreted as a measure of voting power at all. For as Felsenthal and Machover (1998, 195, 275) argue, not only is the additivity axiom used to justify *SS* unmotivated for measures of a priori voting power, *SS* also violates the added blocker postulate, which, they argue, an a priori measure of voting power ought to satisfy. (We argue in Abizadeh and Vetta (working paper-a, working paper-b) that their formulation of the postulate is not reasonable, but that *SS* nevertheless still violates the postulate once reasonably reformulated.) As an a priori measure, *SS* is therefore best interpreted as an index of the relative *value* of a voter's a priori voting power: for example, as the player's expected payoff assuming a cooperative game with transferable utility (Felsenthal and Machover 1998), or as a relative bribe index (Morriss 2002).

Acknowledgements

I am grateful to Adrian Vetta for comments and for jointly developing with me the specific formula for the Recursive Measure. Thanks also to Rohan Dutta, several anonymous referees, and the editor for comments.

References

Abizadeh, Arash. 2021a. Counter-Majoritarian Democracy. *American Political Science Review* 115(3):742-756.

———. 2021b. The Power of Numbers. *Philosophy & Public Affairs* 49(3):290-318.

———. forthcoming. The Grammar of Social Power. *Political Studies*.

Abizadeh, Arash, and Adrian Vetta. Working paper-a. The Blocker Postulates for Measures of Voting Power.

———. Working paper-b. A Recursive Measure of Voting Power that Satisfies Reasonable Postulates. https://arxiv.org/abs/2105.03006

Albert, Max. 2003. The Voting Power Approach. *European Union Politics* 4(3):351-366.

Allen, Amy. 1998. Rethinking Power. *Hypatia* 13(1):21-40.

Amar, Akhil Reed. 1984. Choosing Representatives by Lottery Voting. *Yale Law Journal* 93:1283-1308.

Banzhaf, John F. 1965. Weighted Voting Doesn't Work. *Rutgers Law Review* 19:317-343.

———. 1966. Multi-Member Electoral Districts. *The Yale Law Journal* 75(8):1309-1338.

Barry, Brian. 1988. The Uses of ‘Power’. *Government and Opposition* 23(3):340-353.

———. 2002. Capitalists Rule Ok?. *Politics, Philosophy & Economics* 1(2):155-184.

Braham, Matthew. 2008. Social Power and Social Causation: Towards a Formal Synthesis. In *Power, Freedom, and Voting*, ed. M. Braham and F. Steffen. Berlin: Springer.

Braham, Matthew, and Manfred J. Holler. 2005. The Impossibility of a Preference-Based Power Index. *Journal of Theoretical Politics* 17(1):137-157.

Braham, Matthew, and Martin van Hees. 2009. Degrees of Causation. *Erkenntnis* 71(3):323-344.

Bratman, Michael E. 2014. *Shared Agency*. Oxford: Oxford University Press.

Chant, Sara Rachel, and Zachary Ernst. 2007. Group Intentions as Equilibria. *Philosophical Studies* 133(1):95-109.

Diskin, Abraham, and Moshe Koppel. 2010. Voting Power: An Information Theory Approach. *Social Choice and Welfare* 34(1):105-119.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper & Brothers.

Felsenthal, Dan S., and Moshé Machover. 1998. *The Measurement of Voting Power*. Cheltenham, UK: Edward Elgar.

———. 2003. The Voting Power Approach. *European Union Politics* 4(4):473-79.

———. 2004. A Priori Voting Power: What is it All about? *Political Studies Review* 2(1):1-23.

———. 2009. A Note on Measuring Voters' Responsibility. *Homo Oeconomicus* 26(2):259-271.

Garrett, Geoffrey, and George Tsebelis. 1999. Why Resist the Temptation to Apply Power Indices to the European Union? *Journal of Theoretical Politics* 11(3):291-308.

Goldman, Alvin I. 1972. Toward a Social Theory of Power. *Philosophical Studies* 23(4):221-268.

———. 1974. On the Measurement of Power. *Journal of Philosophy* 71(8):231-252.

Hitchcock, Christopher. 2001. The Intransitivity of Causation Revealed in Equations and Graphs. *The Journal of Philosophy* 98(6):273-299.

Jones, Peter. 1983. Political Equality and Majority Rule. In *The Nature of Political Theory*, ed. D. Miller and L. Siedentop. Oxford: Clarendon Press.

Laruelle, Annick, and Federico Valenciano. 2005a. Assessing Success and Decisiveness in Voting Situations. *Social Choice and Welfare* 24(1):171-197.

———. 2005b. A Critical Reappraisal of Some Voting Power Paradoxes. *Public Choice* 125(1):17-41.

———. 2008. *Voting and Collective Decision-Making*. Cambridge: Cambridge University Press.

Leech, Dennis. 1990. Power Indices and Probabilistic Voting Assumptions. *Public Choice* 66(3):293-299.

Lewis, David. 1986. *Philosophical Papers: Volume II*. Oxford: Oxford University Press.

List, Christian, and Philip Pettit. 2011. *Group Agency*. Oxford: Oxford University Press.

Machover, Moshé. 2007. Voting Power when Voters' Independence is not Assumed. *LSE Research Online: Discussion Topic*. http://eprints.lse.ac.uk/2966

McDermott, Michael. 1995. Redundant Causation. *The British Journal for the Philosophy of Science* 46(4):523-544.

Morriss, Peter. 2002. *Power: A Philosophical Analysis*. 2nd ed. Manchester: Manchester University Press.

Napel, Stefan. 2019. Voting Power. In *The Oxford Handbook of Public Choice, Volume 1*, ed. R. D. Congleton, B. Grofman and S. Voigt. Oxford: Oxford University Press.

Napel, Stefan, and Mika Widgrén. 2005. The Possibility of a Preference-Based Power Index. *Journal of Theoretical Politics* 17(3):377-387.

Owen, Guillermo. 1975. Multilinear Extensions and the Banzhaf Value. *Naval Research Logistics Quarterly* 22(4):741-750.

Penrose, L. S. 1946. The Elementary Statistics of Majority Voting. *Journal of the Royal Statistical Society* 109(1):53-57.

Ramachandran, Murali. 1997. A Counterfactual Analysis of Causation. *Mind* 106(422):263-277.

Saunders, Ben. 2010. Democracy, Political Equality, and Majority Rule. *Ethics* 121(1):148-177.

Schaffer, Jonathan. 2003. Overdetermining Causes. *Philosophical Studies* 114(1):23-45.

Schmidtchen, Dieter, and Bernard Steunenberg. 2014. On the Possibility of a Preference-Based Power Index. In *Voting Power and Procedures*, ed. R. Fara, D. Leech and M. Salles. Cham: Springer International Publishing.

Searle, John R. 2010. *Making the Social World*. Oxford: Oxford University Press.

Shapley, L. S., and Martin Shubik. 1954. A Method for Evaluating the Distribution of Power in a Committee System. *The American Political Science Review* 48(3):787-792.

Straffin, Philip D. 1977. Homogeneity, Independence, and Power Indices. *Public Choice* 30(1):107-118.

Tuck, Richard. 2008. *Free Riding*. Cambridge, MA: Harvard University Press.

Turnovec, František, Jacek Mercik, and Mariusz Mazurkiewicz. 2008. Power Indices Methodology. In *Power, Freedom, and Voting*, ed. M. Braham and F. Steffen. Berlin: Springer.

Weber, Max. 1978. *Economy and Society*. Berkeley, CA: University of California Press.

Weber, Robert James. 1988. Probabilistic Values for Games. In *The Shapley Value*, ed. A. E. Roth. Cambridge: Cambridge University Press.

Wright, Richard W. 1985. Causation in Tort Law. *California Law Review* 73:1735-1828.

———. 1988. Causation, Responsibility, Risk, Probability, Naked Statistics, and Proof. *Iowa Law Review* 73(1001-77).

[1] For the measure's initial formulations, see Penrose (1946); Banzhaf (1965, 1966).

[2] For a defence, see Tuck (2008, 40-41). Braham and van Hees (2009, 333) take the opposite view; for criticism, see Felsenthal and Machover (2009).

[3] That is: the unconditional probability any v_i votes for an alternative equals the probability she does so conditional on a particular vote configuration of all other voters; and the unconditional probability of any vote configuration of all other voters equals its probability conditional on v_i voting for a given alternative.

[4] Albert (2003) deems the principle discredited; Felsenthal and Machover (2003, 478; 2004, 15) reply that it is unproblematic for a finite probability space comprising finite atomic events.

[5] *GPB* is thus equivalent to Φ in Laruelle and Valenciano (2005b, 2008).

[6] But see Napel and Widgrén (2005); Schmidtchen and Steunenberg (2014). Because it fails to counterfactualize over the power-assessee's preferences, Diskin & Koppel's (2010) proposed measure violates this condition and hence cannot be interpreted as a measure of agential power.

[7] See Allen (1998); Abizadeh (forthcoming) on power-with.

[8] For related approaches, see McDermott (1995); Ramachandran (1997).

[9] Cf. the similar approach in Braham and van Hees (2009).

[10] There are two differences. First, Braham and van Hees calculate degrees by counting a division itself as one its own loyal descendants (i.e., counting *subsets*, not just *proper* subsets, of the actual SSA). Second, Braham and van Hees (2009, 333) assume (contra both Tuck and *PB*) that each actor's degree of causal contribution represents a *portion* of a constant sum of total causal efficacy. For criticism, see Felsenthal and Machover (2009). Braham and van Hees, moreover, impose this constant-sum efficacy assumption inconsistently: on actors' efficacy in the parent *division*, but not *loyal descendants*. [E.4] imposes the non-constant-sum assumption consistently. See also

Goldman (1972, 241; 1974, 239).

[11] The specific formula of the Recursive Measure's division efficacy score (and the loyal child concept on which it is based) was developed jointly with Adrian Vetta, and is justified by the lattice representation of voting games and the notion of random walks in stochastic processes. See Abizadeh and Vetta (working paper-b).

[12] Why do we weight a loyal child's division efficacy scores by the number of loyal children, rather than its probability given the distribution of preferences (and strategic interaction)? Because unlike when we calculate *voting power* on the basis of division efficacy scores, where the weights attaching to those scores may represent probabilities, when we calculate *efficacy scores* themselves the weights attaching to their loyal children do not admit of a probabilistic interpretation. Whereas in calculating voting power we adopt the ex ante perspective prior to an actual division (and so interpret division weights as probabilities), in calculating efficacy scores we adopt an ex post perspective from which the division (whose efficacy score is being calculated) has already occurred. From this ex post perspective, the probability of any descendant is zero; we are instead considering the *proportion* of loyal children in which the voter would be efficacious, and the *degree* of efficacy were the loyal child to occur, not the probability it might occur.

[13] Representing divisions as ordered lists of actions, rather than sets of indexed action-events.

[14] E.g. the loyal grandchildren of (y,y,y,y,y) begotten by (y,y,y,y,n) are: (y,y,y,n,n), (y,y,n,y,n), (y,n,y,y,n), (n,y,y,y,n).