

ENDOGENOUS ORGANIZATIONAL RESTRUCTURING: STATUS, PRODUCTIVITY & MERITOCRATIC DYNAMICS

ASHUTOSH THAKUR*

National University of Singapore
Singapore, 259772
adthakur@nus.edu.sg

JONATHAN BENDOR

Stanford Graduate School of Business
Stanford, CA, 94305, USA
jonathan.bendor@stanford.edu

April 20, 2022

ABSTRACT. We model the dynamics of endogenous organizational restructuring where those being assigned positions in an organization can themselves lobby for who gets which position. Internal labor market changes depend on how much individuals value their personal status, organizational output, their friends' welfare, and the quality of their departmental colleagues. We show that an organization converges to the meritocratic, efficient assignment of people to positions by a combination of agents valuing organizational output and restrictions on the scale of reorganization. However, concentrated decision-making power, lax restrictions on agenda-setting protocols, expansive friendship networks, and large departments can hinder such paths of restructuring.

* Corresponding Author is Ashutosh Thakur, National University of Singapore.
Email: adthakur@nus.edu.sg, *Telephone:* 1-520-465-6353

A special thanks to Steven Callander, Alvin Roth, Paulo Somaini, Ken Shotts, and participants of the Political Economy Seminar at Stanford and the Behavioral Management Science Workshop at University of Cologne. Ashutosh Thakur would like to thank Stanford Graduate School of Business and ECONtribute: Markets and Public Policy at University of Cologne, and acknowledge support from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1–390838866.

1. INTRODUCTION

Internal politics in an organization is often complicated because it is an arena where collective goals and private interests can collide. For example, members often benefit indirectly from improved organizational performance but this may be at odds with their personal ambitions to seek better positions for themselves within the organization despite others being more qualified for the job.

Since people care a great deal about their own status, an especially contentious issue is who gets what position within the organization. And because many organizations have considerable discretion over their own personnel decisions, such issues arise often.

This gives scope for coalitions to arise endogenously and lobby for internal reorganizations that favor themselves. Which coalitions arise, what they push for, and whether they are successful in bringing about their intended reorganizations depends on individual's *preferences* over positions, the institutional *procedures* restricting the scope of the reshuffling, and the group's *power* to institute change.

Whether it is politicians being assigned to committees in a legislature, workers being assigned to positions in a firm, bureaucrats being assigned to jobs in the civil service, or cadets being assigned to various branches of the army, internal labor market changes can impact the satisfaction of the organization's members, the cohesiveness of its teams, the longevity of its departments, and the efficiency of its production.¹ How do organizations endogenously evolve when their members push for changes that boost their own status, the organization's output, and the well-being of their friends and of their departments? Which allocations are stable, are they efficient and meritocratic in their assignment of talent? What is the organization's path of restructuring?

To study these dynamics of organizational restructuring, we apply tools from matching theory and social choice.

¹See Thakur (2017, 2019), Cowgill and Koning (2018), and Sonmez and Switzer (2013).

Assigning the N members of an organization to N different positions over which they have *preferences* is a one-to-one matching problem at its core. However, our paper is not based on some underlying centralized matching mechanism.² Instead, we study how an organization's path of personnel reshuffling is driven by the endogenous pressure imposed by its members.

Which coalitions succeed in replacing the status quo allocation with their favored allocation in every period often depends on a complex mix of informal as well as formal procedures. We represent this as a formal process using tools from social choice: voting power and procedural constraints on the agenda. Voting is an idealization of often informal pressure on an organization to alter its status hierarchy or otherwise restructure, although sometimes formal voting is used.³ The organization's *power* structure determines who has how much say and hence which coalitions are strong enough to bring about reorganizations. Lastly, *procedural* constraints define the scope of reorganization that is entertained in every period: how the set of alternatives that can challenge the status quo allocation is constrained. The combination of agents' preferences, voting power, and agenda-setting restrictions establish a well-defined dynamic that we can study.

We analyze these issues by constructing a dynamic model of person-position matching inside an organization. In every period, N agents are assigned to N seats, a one-to-one matching. In the baseline model, positions are ordered hierarchically by status $1, \dots, N$. Each agent has a quality level and agents can be ordered by skill $1, \dots, N$. Each assignment of agents to seats produces a level of output; assigning higher quality agents to better seats generates higher output. Agents get utility from the status of their assigned seat, the level of output produced by the overall assignment, the status of their friends, and the quality of agents assigned to their team. The alternative allocation protocol defines the set of alternative allocations that can challenge the status quo. The status quo allocation is replaced by the alternative if it passes the

²Thakur (2020) studies institutional stability, treating an organization's matching mechanism as an institutional choice.

³For example, when political party members vote on changes to committee assignments and assignment procedures in a legislature or when shareholders with voting rights propose structural changes.

voting rule used by the organization. We analyze the set of stable allocations which cannot be overturned by any alternative allocation and the dynamic — the organization’s path of restructuring — by which they might be reached.

Our main finding, Theorem 1, is that if agents even slightly value organizational output then organizations move to and stay at the first best allocation if and only if there are institutional restrictions on the number of promotions allowed under the alternative allocation. In this stable and efficient allocation agents are assigned meritocratically based on quality. The convergence to and absorption into the first best allocation holds even when there is maximal rivalry amongst agents due to their preferences over seats being fully correlated (e.g., a commonly held status hierarchy).

A crucial aspect of the dynamics is that for individuals whose positions are not changed by a proposed alternative allocation (“*uninvolved agents*”) and who thereby might be indifferent to the proposed change, valuing organizational output even very slightly induces them to support output-improving changes and oppose those that would reduce output. Thus, these otherwise indifferent agents join pro-efficiency coalitions and support reorganizations that increase output (Corollary 1.1) and enhance meritocracy (though not necessarily monotonically as Remark 3 shows).

This disciplining power of the uninvolved agents is significant: Proposition 1 shows this force holds even for non-hierarchical organizations where agents have varied preferences over their positions. However, attaining the unique, stable, and efficient outcome requires more than agents who value organizational output. It is also necessary to impose procedural restrictions on the number of promotions allowed under the alternative allocation. Otherwise, the endemic rivalry for positions amongst agents in hierarchical organizations can produce instability and chaos, even if all agents also value organizational output. We highlight this in Proposition 3 where we derive a chaos theorem akin to McKelvey (1976): if the alternative allocation protocol is unrestricted, then even when agents prefer more output to less, there exists a majority-approved sequence of reassignments from any allocation to any other allocation. Restricting the scope of organizational shake-ups does imply the possibility of organizational inertia, i.e., reorganization efforts that fail, thereby retaining

the status quo. Of course, failed reorganization efforts and the ensuing inertia are empirically rather common, so it is important that they are implications of a descriptive theory of intra-organizational conflict over positions. But we also show that these phenomena are part of what helps discipline a potentially chaotic process where agenda setters would otherwise enjoy arbitrary power.

Nevertheless, we also show that even when agenda setting is unconstrained, the first best allocation has descriptive significance. Theorem 2 establishes, under plausible assumptions about the probabilities that different agents will be recognized as agenda setters, that the first best holds a privileged position in the limiting distribution of the underlying Markov process.

We then examine organizational features that can impede the movement towards the first best. We find that both formal and informal organizational structures (departments and friend networks, respectively) as well as concentrated decision-making power can offset the forces pushing for pro-efficiency reorganizations. The richer organizational structures result in there being fewer uninvolved agents since friendship and departments increase agents' skin in the game. Concentrated decision-making power decreases the ability of uninvolved agents to express their voice.

In Section 4, we illustrate organizational impediments to reaching efficiency arising from the network of informal ties and connections and from the formal team/department structure commonly found in many organizations. To restore stability of the first best, the alternative allocation protocol has to be further restricted. In the absence of stringent restrictions on the alternative allocation protocol, even when agents value organizational output, unmeritocratically promoted agents and their friends can form sizeable coalitions even with sparse friendship networks (Remark 4) and departments can act as voting blocs who benefit from selfishly boosting their own teams' talent pool (Remark 5). Even when the alternative allocation protocol is most severely restricted to allowing only one promotion at a time, if an agent is extremely well-connected or a friend clique or department is too large, the first best is not stable (Propositions 4 and 5). Our model then illustrates why in a family firm a kin-based coalition might block the promotion of talented outsiders into upper

level management and retain less productive dynastic family CEOs (Perez-Gonzalez, 2006). Relatedly, disproportionately large coalitions of cliques or high-solidarity groups based on race or gender can block certain assignments and generate segregation, discrimination, and glass-ceiling effects within the workplace (Smith, 2002). Lastly, powerful large departments can impair the functioning of smaller units: e.g., Xerox’s copier division dominated Xerox’s PARC division thereby the company failed to monetize its innovation of the earliest personal computer. Crucially, our results underscore the relation between coalitional politics in organizations and the scope of the reorganization admitted by the alternative allocation protocol.

In Section 5, we consider many voting rules an organization may use to resolve internal conflict. In particular, we examine dictatorship, oligarchy, veto power, and other rules that allow certain agents to dominate decision making. Such rules involve powerful agents who are part of every decisive coalition — part of a “collegium” as it is called in the social choice literature. These pivotal decision makers can ensure that the meritocratic discipline desired by uninvolved agents is not imposed on them, but on everyone else! Theorems 3 and 4 generalize Theorem 1 by considering more informal processes of influence, persuasion, negotiation and debate whereby agents’ power can depend arbitrarily on who is assigned to which position.

Lastly, in Section 6 we relax the assumption that agents vote myopically between the status quo and an alternative allocation in each period. We consider more sophisticated voting by analyzing forward-looking opportunity costs, combinations of short-run and long-run considerations, and deploying the solution concept of uncovered sets. Proposition 6 and Theorem 5 show that a central property of our model — uninvolved agents having even an epsilon preference for output powerfully bends the path of restructuring toward the first best — is robust against many kinds of sophisticated voting.

Related Literature. Whereas Thakur (2020) considers a static problem—the existence of institutionally stable allocations—our paper analyzes an organization’s dynamic path of restructuring. Further, in addition to the preference assumed by Thakur (rank order preferences over one’s assigned seat), we allow

for preferences for organizational output, friendship networks, and the quality of departments.

There is a small yet growing literature on popular matching (i.e., stability under plurality rule)—originating from Abraham et al., (2007)—and its many variants in the computer science literature (see review articles by Manlove [2013] and Cseh [2017]). This research focuses on the the existence, characterization, and complexity of finding a matching allocation that is a Condorcet winner under plurality rule given one-sided voting, two-sided voting in bipartite graphs, and non-bipartite graphs. Although the focus on algorithmic complexity has made this work normatively appealing, it is not linked to applications nor has it received much attention from the literature on economics and applied market design. In addition to including rank order preferences, we enrich the model with several empirically common motives that affect agents’ preferences over positions in organizations, bringing our model closer to real-world contexts. Furthermore, we study the dynamic path of matching allocations reached endogenously under various agenda restrictions and many different voting rules, not just plurality rule.

Lastly, analyzing the stability of the institutional choice itself shares the spirit of self-stable voting rules and the constitutional choice literature (Barbera and Jackson, 2004; Messner and Polborn, 2004; Maggi and Morelli, 2006; Koray, 2000). However, in our model agents do not choose a voting rule; instead, they collectively and dynamically choose matching allocations for themselves.

2. MODEL

Consider the problem of assigning $N \geq 3$ **agents** (denoted by $i \in \mathbb{I} = \{1, \dots, N\}$) to N **seats** (denoted by $s \in \mathbb{S} = \{1, \dots, N\}$) within an organization. For ease of notation, assume N is odd, unless otherwise stated. An **allocation** A is a one-to-one matching between agents and seats.

Agent i receives **status** utility $u_i(s)$ from being assigned to seat s . We assume that there is a perfectly agreed-upon status hierarchy over seats, i.e.,

all agents have identical *ordinal* preferences over seats⁴ in that $u_i(s) > u_i(s')$ if $s < s'$ for all i .⁵ This is relaxed in Section 3.1.

Each agent i has a **quality** or skill that is decreasing in i ; hence let i denote a player's type, where a lower-numbered type is a more skilled or higher quality agent. An allocation A produces an **output** $\mathbb{Y}(A)$. Let $A(i)$ denote the seat that i is assigned in allocation A . We assume that the output production function $\mathbb{Y}(\cdot)$ satisfies the following condition: if $i < j$ but $A(i) > A(j)$ then an alternative allocation A' that swaps i and j , while leaving all other assignments unchanged, increases output: $\mathbb{Y}(A') > \mathbb{Y}(A)$. This means that the marginal value of having a higher quality agent is more in a higher status seat.⁶

We model two other universal features of organizations: friend networks/connections and teams/departments.

Friends, relatives, professional acquaintances, and other personal relationships all lead to complex networks of informal ties within an organization that can affect career opportunities and internal politics. We assume that agent i has a set of **friends**, $Friend_i \subseteq \mathbb{I} \setminus i$, who he cares about and derives utility from their status. We do not require that ties be symmetric: i can consider j his friend but j might not reciprocate. Asymmetric ties could arise for several reasons: e.g., although both i and j could have friendly feelings toward each other, j likes to insulate the workplace from personal ties whereas i is comfortable being friends with co-workers. Although connections (especially those made for career advancement) are endogenously and strategically formed, analyzing this complex phenomenon is a major topic in its own right; hence we assume that the friendship network in the organization is exogenously specified.

⁴From Thakur (2020) we know that this assumption of perfectly correlated preferences is the worst case for organizational stability because it implies maximal rivalry across agents.

⁵We assume that the salaries corresponding to the positions are already reflected in the perfectly agreed-upon status hierarchy. The cardinal utility difference between two seats can also capture endogenous bonuses/salary differences (e.g., arising from stack ranking systems or rank-order contests [Lazear and Rosen, 1981]) between two otherwise equivalent positions at the same level or seniority.

⁶Increasing differences over type and position for output is a sub-case of this restriction on the output production function: for type $i < j$ and seat $s < s'$, $\mathbb{Y}(A(i) = s) - \mathbb{Y}(A(j) = s) > \mathbb{Y}(A(i) = s') - \mathbb{Y}(A(j) = s')$ holding constant all other agents' assignments in A .

Most organizations of even modest size have some team/departmental structure. Naturally, agents care about who else is in their team or department. Having more capable individuals assigned to one's team can yield numerous benefits: group incentive pay, improving division of labor, increasing group expertise, and improving work satisfaction due to the increased competence of colleagues. We assume that the set of seats is partitioned into teams, and each agent i belongs to exactly one **team**, $Team_i \in \mathbb{S}$ depending on the seat he is assigned.

The **utility of agent i from allocation A** , $U_i(A)$, is defined by

$$(1) \quad U_i(A) = \alpha u_i(A(i)) + \beta \mathbb{Y}(A) + \gamma \sum_{j \in Friend_i} u_j(A(j)) + \zeta \frac{1}{|Team_i|} \sum_{s \in Team_i} f_i(s(A))$$

where $s(A)$ is the agent who is assigned seat s in allocation A , and f_i are decreasing functions (i.e., the agent benefits from having teammates with lower index, corresponding to higher skill).

The first term is the agent's payoff from the status of his position. A person has preferences over being assigned to different positions and hence may be envious of other members' positions in the organization. These may arise from a desire for career advancement, better work-life balance, or increased power or responsibility.

The second term is the payoff from the performance of the organization as a whole. Workers may have stock-based compensation directly reflecting their company's productivity (Hannes, 2006), politicians might want to pass legislation to claim credit and boost their re-election chances (Mayhew, 1974), and an army's combat effectiveness can be a matter of life or death for a cadet.

The third term is the payoff that agent i derives from the status enjoyed by i 's friends given their assigned seats. This can account for kin, professional contact, or any other social tie. Members' social ties can induce preferences over what positions one's friends, family members, and colleagues get. Politicians belonging to the same political faction (Francois et al., 2016), bureaucrats with hometown or college connections (Fisman et al., 2020), and workers of the same gender, race, or ethnicity (Smith, 2002) might support the ambitions of those in their in-group.

The fourth term is the payoff from the quality of those assigned to your team. Having more capable individuals on one’s own team can enhance group-expertise — e.g., politicians with expertise and experience in the domain of their legislative committee help make better-informed choices (Krehbiel, 1992) — or increase bonuses from incentive schemes tied to team output (Boning et al., 2007).⁷

We are interested in the dynamics of how an organization is restructured over time. There are discrete time periods $t = 0, 1, \dots$ and the corresponding allocations of agents to seats in each period are $A_0, A_1, \dots, A_t, \dots$

In each period t an agent i **votes** over two allocations if and only if he has a strict preference between them; otherwise he abstains. The voting rule for the organization’s allocation choice is majority rule: alternative allocation A'_{t+1} replaces A_t if and only if a majority of agents vote for A'_{t+1} over A_{t+1} . We consider other common voting rules in Section 5 and more general analysis using decisive coalition characterizations in Section 3.3.

We assume myopic voting, but relax this to incorporate some forward-looking, sophisticated voting in Section 6.

The assumption of voting is of course sometimes an idealization of the often informal pressure on the organization to undergo restructuring and reorganization. When group decisions are reached organically through debate, discussion, and deliberation, we interpret voting to be an idealization of providing an opinion in favor of or against a proposed change. Some organizations, however, literally use formal voting when restructuring, e.g., when political party members vote on alternative committee assignment procedures and change allocations in the legislature⁸ or when a company’s employees who are compensated via stocks with voting rights propose structural and personnel changes.⁹ Further, subjective peer-evaluations in the form of performance appraisal reports or 360 degree evaluations from managers, peers, and subordinates often

⁷With added assumptions this could also be written in terms of output produced by those who are assigned to your team, e.g., by assuming that output is additive across departments/seats. This latter assumption would enable us to isolate each department’s specific contribution to overall organizational output. However, to impose less structure (e.g., to allow for more general complementarities across departments), we use the fourth term to represent quality, which in our model is intrinsically tied to output.

⁸E.g., in anticipation of the impeachment hearings for President Trump, “some GOP mem-

help to determine promotions. Moreover, even in debates and deliberations like-minded agents mobilize and form coalitions advocating for changes that work out in their favor.

We also consider voting rules other than majority rule to capture inequalities in power: the important property that the opinions of certain agents or of agents in certain positions can be more influential than others. For example, in determining a party's committee assignments, the senior party leadership and ranking party members may have greater voice, the Board of Directors or a CEO may have veto power to end debate over proposals raised by a firm's lower-level employees, or a firm considering promoting its employee may ask for performance appraisals only from that employee's managers and workers working directly under that employee's supervision. Organizations discipline their internal decision-making in different ways; examining a set of diverse voting rules helps us capture some of this empirical richness, albeit in a stylized manner.

In each period t , the set of alternative allocations that can challenge the status quo A_t is determined by the *alternative allocation protocol*. We say that the “*alternative allocation protocol is restricted to some property*” to signify that the admissible set of alternative allocations that can challenge any status quo in any time period t must satisfy the property. Unless otherwise stated, we assume **Assumption A**: there exists an $\epsilon > 0$ such that the probability that an admissible allocation A is proposed to challenge the status quo in any period t is at least ϵ . Our results do not depend on exact specifications of the agenda over alternative allocations and even this

bers have expressed frustration that some Republicans on the House Intelligence Committee have not been participating in all the depositions, unlike Jordan and some of Trump's other top allies.” The Republican party leadership thus announced a reshuffling of committee assignments: “House Minority Leader Kevin McCarthy (R-Calif.) announced the roster shakeup on Friday, adding that Rep. Rick Crawford (R-Ark.) has agreed to temporarily step aside for Jordan (R-Ohio), who is seen as one of Trump's best attack dogs on Capitol Hill.” <https://www.politico.com/news/2019/11/08/jim-jordan-intelligence-committee-068049>

⁹E.g., consider the employee activism spurred by significant share-holding with voting rights of Amazon employees via their compensation. “Amazon employees are using their company-issued stock to pressure top executives into reducing contributions to climate change,[...], transparency regarding the lack of ideological diversity on the board...” <https://www.wired.com/story/amazon-employees-try-new-activism-shareholders/>

structure of probabilistic alternatives challenging the status quo can be further relaxed, though we maintain it for ease of exposition. Note that Assumption A does not require that the alternative allocation protocol be stationary (e.g., allocations often rejected in the past could become less likely to be proposed subsequently) or that it be Markovian (what is proposed today may depend on what happened along the path that led up to the present state). We study the full range of restrictiveness on the alternative allocation protocol, from allowing any alternative allocation to the maximal restriction of only allowing pairwise swaps.

Together the voting rule and the alternative allocation protocol determine the ***path of restructuring***, $\{A_t\}_{t>0}$, which refers to the consecutive sequence of allocations reached. This includes consecutive instances of the same allocation in periods $t - 1$ and t arising from an alternative allocation being voted down. An allocation A_t^* is said to be ***stable under a voting rule*** for some alternative allocation protocol if there is no admissible alternative allocation that can defeat allocation A^* under the voting rule and hence $A_s = A_t^*$ for all $s \geq t$. When using the terms “***stable***” or “***stability***” without qualifying any voting rule, we are referring to what is stable under majority rule.

In this paper we analyze the path of restructuring traversed by the organization in response to how much agents value different organizational features (the utility weights), what kinds of reshuffling are considered (the alternative allocation protocol), and how internal conflicts are resolved (the voting rule).

3. THE CONSEQUENCES OF VALUING ORGANIZATIONAL OUTPUT.

We start the analysis by studying agents whose utility depends only on their own seat assignment. We then consider the incremental effect of agents also valuing organizational output.

Remark 1. *If $\alpha = 1$ and the alternative allocation protocol is unrestricted then no allocation is stable.*

The shared status ranking of positions in the organization implies that preferences are fully correlated. Hence, no allocation is stable under majority rule as $N - 1$ agents can always be made better off at the expense of the agent assigned the top seat. This instability is robust to the collective choice procedure:

it holds for any q -rule that is not unanimity rule. This shows how powerful the unrestricted alternative allocation protocol is at destabilizing status quo allocations, given perfectly correlated preferences over seats and the ensuing rivalrous nature of the private-status-only environment.

Stability under majority rule can be restored in this environment if the alternative allocations proposed satisfy certain criteria. For example, suppose that alternative allocations are restricted to a pairwise swap: i takes j 's seat and vice versa. This restriction makes *every* allocation stable; only i will vote for her proposal; j will vote against it and the $N - 2$ uninvolved agents, being indifferent, will abstain, thus killing the proposal. Thus restricting to pairwise swaps restores majority stability, but at the price of extreme institutional inertia: even grossly inefficient allocations of people to positions cannot be changed.

The restriction to pairwise swaps is only an example of an alternative allocation protocol that produces extreme stability in the private-status-only environment. A large class of protocols has this effect.

Let N_p be the number of promotions (i.e., agents improving status) proposed by the alternative allocation, N_d the number of demotions (i.e., agents losing status), and N_u the number of *uninvolved agents* whose positions remain unchanged. For any alternative allocation proposed relative to the status quo, $N_p + N_d + N_u = N$.

Remark 2. *If $\alpha = 1$ and the alternative allocation protocol is restricted to those in which $N_p \leq N_d$ then every allocation is stable under majority rule.*

Thus a large class of protocols restore stability under majority rule, but at the price of making *everything* stable: the opposite problem of extreme instability.

It turns out, however, that the problem of extreme stability reflects a knife-edge condition. Suppose that the agents care about output, though perhaps only slightly: β is some value $\epsilon > 0$, while $\alpha = 1 - \epsilon$. Because the status of the uninvolved agents is not at stake, they care only about organizational performance; hence all N_u of them will vote for whichever allocation produces more output. Hence, if the status quo is maximally efficient then the uninvolved agents will vote against any challenger.

This amounts to a proof for the following useful lemma.

Lemma 1. *If $\alpha + \beta = 1$, $\beta > 0$, and the alternative allocation protocol is restricted to those in which $N_u \geq \frac{N+1}{2}$ then the first best allocation is stable under majority rule.*

This restriction is justified as organizational policy because the agents whose status is unaffected by the vote can be expected to focus their attention on organizational performance. This is an intuitive property that is captured by the preference-assumptions of Lemma 1 and our model's basic utility assumption. Hence, if these (temporarily neutral) arbiters of a status-fight are a majority in the organization then their views will prevail, thus protecting the status quo if it is efficient.

Indeed, the assumptions of Lemma 1 are stronger than needed. In order to stabilize an efficient status quo it is not necessary that the uninvolved agents be a majority. It is only required that those who would attack such a status quo be a minority in the organization. Therefore, henceforward we will mostly use the following more general lemma.

Lemma 2. *If $\alpha + \beta = 1$, $\beta > 0$, and the alternative allocation protocol is restricted to those in which $N_p < \frac{N+1}{2}$ then the first best allocation is stable under majority rule.*

Moreover, this does not recreate the above problem of extreme and dysfunctional organizational inertia: given even an ϵ -weight on output, every inefficient status quo allocation is unstable.

Lemma 3. *If the assumptions of Lemma 2 hold then all inefficient allocations are unstable.*

The proof follows immediately from two facts: first, the uninvolved agents are a majority in challenges to the status quo which involve only a pairwise swap; second, every inefficient allocation has at least one meritocratically twisted pair (i.e., $i < j$ but $A(i) > A(j)$).

These two lemmas together imply that if the assumptions of Lemma 2 hold then the first best allocation is the uniquely stable allocation.

Finally, the threshold identified by Lemma 2, $N_p < \frac{N+1}{2}$, is necessary as well as sufficient: if it is not satisfied then the first best allocation is unstable.

Lemma 4. *Suppose $\alpha + \beta = 1$ and $\beta > 0$. Then the first best allocation is stable only if the alternative allocation protocol is restricted to those in which $N_p < \frac{N+1}{2}$.*

The proof is simple. Consider a challenge to the first best allocation for which $N_p \geq \frac{N+1}{2}$. Since the status quo allocation is the first best it produces more output than that generated by the new allocation, but for β sufficiently close to zero the agents who would be promoted value their own status improvement more than the output that would be lost. Moreover, they constitute a majority. Hence, they will vote to overturn the status quo allocation.

We now turn to examining the organization's path of restructuring. To do this we study the dynamic process which governs how the organization's allocations evolve. The states of this stochastic process are the allocations. In this process a state is absorbing if and only if the allocation is, in terms of the preceding static analysis, stable. Therefore, if the assumptions of Lemma 2 hold then the stochastic process has a unique absorbing state: the first best allocation.

By itself the existence of a unique absorbing state does not, of course, imply that the process converges to that state. The process might get hung up somewhere: there might be a subset of inefficient allocations that collectively form an absorbing class of states. We show that this would necessitate certain restrictions on what kinds of challenges to the status quo can be proposed: e.g., for some reason new allocations that would defeat an inefficient status quo simply cannot be proposed.

Our main result shows that provided that the assumptions of Lemma 2 hold, the organization's path of restructuring has a clear trajectory.

Theorem 1. *Suppose $\alpha + \beta = 1$ and $\beta > 0$. Then the path of restructuring is absorbed into the first best—the unique stable allocation—with probability 1 if and only if the alternative allocation protocol is restricted to those in which $N_p < \frac{N+1}{2}$.*

Proof: Lemmas 2 - 4 imply that part the first best is stable if and only if the alternative allocation protocol is restricted to those in which $N_p < \frac{N+1}{2}$. Assumption A implies that output-improving pairwise swaps are generated with a probability of at least $\epsilon > 0$ for every inefficient status quo in any period t . Any such swap is approved by a vote of $N - 1$ to one. This means, given that the stochastic process has finitely many states, that there is a finite sequence of steps from any arbitrary inefficient status quo to the first best allocation where each step is realized with positive probability. Hence the first best allocation is reachable from every other state of the process, and because the first best allocation is an absorbing state, once the path of restructuring reaches that allocation it stays at that state thereafter.

An alternative allocation protocol that allows for more promotions can of course provide quicker meritocratic corrections in some well warranted instances. Suppose, for example, that the worst quality agent occupies seat 1 while everyone else is ordered meritocratically. Then the first best allocation can be achieved in one vote — but there is only one way to accomplish that: by demoting the worst agent from seat 1 to N while moving every other agent up one seat. This requires promoting $N - 1$ agents in one fell swoop, which is precluded by the protocol-criterion assumed in Theorem 1.

Nevertheless, that criterion has a strong justification: not only is it a sufficient condition for the stability of the first best allocation, it is also *necessary*. Any criterion that permits more promotions (i.e., $N_p \geq \frac{N+1}{2}$) allows the maximally efficient allocation to be overthrown by a coalition of less-able agents who are willing to sacrifice organizational performance in order to advance their personal interests by securing higher positions. Indeed, allowing for new proposals in which a majority of agents are promoted creates pervasive instability caused by rotating coalitions of agents cycling around seats simply in order to improve their private status-payoff.

When agents also value organizational output, the voting behavior of the *uninvolved agents*—those whose own assignment does not change and who consequently vote only for output-improving proposals—is crucial.

As the rest of the paper will demonstrate, the overwhelming impact of the otherwise indifferent uninvolved agents, their involvement triggered by $\beta > 0$, is a robust effect. Giving employees and executives stock-based compensation has been hotly debated yet frequently used in practice with the intention of aligning incentives and inducing effort.¹⁰ Theorem 1 reveals the effects of such policies on the endogenous pressure for reorganizing personnel and talent within an organization.

The next few subsections will develop the more subtle and nuanced intuition as to how $\beta > 0$ and alternative allocation protocol restrictions go hand-in-hand in generating the (relatively) well-behaved path of restructuring in Theorem 1.

3.1. Path of Restructuring & Organizational Inertia.

An alternative interpretation of $\beta > 0$ is that of making agents identify as “insiders” in the organization, as studied in identity economics, sociology, and management literature (Akerlof and Kranton, 2005). Our model’s dynamics illustrate how insiders work to further the organization’s goal through restructuring and why identity and culture can matter for organizations.

Corollary 1.1. *If the assumptions of Theorem 1 hold then output is non-decreasing over the path of restructuring.*

Proof: in Online Appendix A.

This dynamic property has an interesting implication in settings where there are shocks to the organization’s environment. Such shocks can alter the quality ordering over the agents’ skills: e.g., an innovation can devalue certain competencies. Naturally, a shock at period t can reduce organizational output. From that time forward, however, the corollary implies that output will once

¹⁰Weitzman (1984) argues employee compensation tied to output production can motivate employees better than flat, fixed wages, and this can add up to having substantial macroeconomic consequences. Hannes (2006) argues that providing stock-based compensation to employees can improve the monitoring of other employees. A theoretical literature beginning with Jensen and Meckling (1976) suggests that option contracts can align executives’ incentives with shareholders and induce effort. Empirical evidence is split, though Demsetz and Lehn (1985), Himmelberg et al. (1999), Core and Guay (1999), and Rajgopal and Shevlin (2002) find that granting stock options is consistent with firm value maximization.

again start increasing. The logic of the N_u uninvolved agents simply restarts hill-climbing on output gradients that have been reset by the recent shock.

Note, however, that output is not strictly increasing over the path of restructuring; there may be significant *organizational inertia*. The path of restructuring will remain stagnant if a proposed reorganization is voted down at some period. Failed attempts at overhaul and reorganization are common in practice and in the context of this model. Moreover, although admitting larger reorganizations by relaxing the restrictions over the alternative allocation protocol would allow faster corrections away from inefficient allocations towards efficiency, this reduction in inertia would come at a price: as we explained after Theorem 1, the consequence of such relaxations is systemic instability.

In spite of the organizational inertia, the strong drift towards efficiency under $\beta > 0$ and restricted alternative allocation protocols is quite robust. In fact, although we focus in this paper on the maximally rivalrous setting of a status hierarchy where agents agree that higher positions are more desirable, the dynamic pull towards efficiency does not depend on this. For example, agents could have different preferences over seats¹¹: some people might not want to become CEO of a firm; becoming the President of a university or chair of a department may not be a work-life balance that is for everyone.

Proposition 1. *For any arbitrary preferences of agents over seats, suppose $\alpha + \beta = 1$ and $\beta > 0$. Then the path of restructuring is absorbed into the first-best—the unique stable allocation—with probability 1 if and only if the alternative allocation protocol is restricted to those in which $N_u \geq \frac{N+1}{2}$.*

Proof: in Online Appendix A.

Proposition 1 underscores the disciplining power of the uninvolved agents voting en bloc only for output-improving reorganizations. The result shows that when agents' preferences over seats may be varied, this disciplining power is not only sufficient, but also *necessary* for stability of the output-maximizing allocation. If agents have heterogeneous desires — some people do not want

¹¹When the alternative allocation protocol is unrestricted and $\alpha = 1$, Thakur (2020) finds that whether an allocation is stable depends on how correlated preferences are and whether preferences admit chains of envy where i takes j 's seat who takes k 's seat and so on, thus unraveling majority stability.

to climb the organizational ladder — restrictions on the alternative allocation protocol must be tighter than they are in Theorem 1. In such circumstances it is possible, per the above proof regarding necessity, for *everybody* whose position would be changed by a new allocation to become better off. This is impossible in the completely rivalrous world of Theorem 1. Although the strategic problem — the tension between collective performance and individual self-interest — is essentially the same in the more complex world of Proposition 1, it is just not coordinated in a straightforward way. Thus, because the uninvolved agents are the only ones who will focus single mindedly on output, they must constitute a majority if the dynamic is to have the efficiency properties identified by Theorem 1.

The organizational drift towards efficiency has an important unintended consequence highlighted by Proposition 1. Under the first best, the preferences of many agents over seats may be ignored: the uninvolved agents might force an involved agent into a position s/he despises. Hence the allocation that maximizes output might differ from the social planner’s first best, which would maximize the (weighted)-sum of everyone’s utilities. The social planner’s first best balances output and positional preferences, but the uninvolved agents can push for efficiency so vigorously that it can take the organization to an extreme that optimizes only on output! This highlights the downside of giving agents stocks/shares of the organization (Weitzman, 1984).

Agents valuing organizational output with $\beta > 0$, however, does *not* imply that allocations are necessarily getting more meritocratic over the path of restructuring. The underlying intuition is that since the number of agents being reshuffled to different positions is allowed to be large, some un-meritocratic moves can be piggybacked onto other personnel changes that boost output enough to win the support of the uninvolved agents.

Remark 3. *If the assumptions of Theorem 1 hold then the distance from the meritocratic ideal, $\sum_i |s_i - s_{\text{meritocraticideal}}|$, can increase in certain periods of the path of restructuring. See Appendix A for an example.*

3.2. Agenda Setting Power & Other Alternative Allocation Protocols.

The restriction of an agenda setter’s power to manipulate the allocation (only output-improving paths of restructuring are majority-approved) comes from *both* the set of admissible alternative allocations being restricted *and* $\beta > 0$. We already know about the properties of one extreme case: restricting the alternative allocation protocol to swaps with $\alpha = 1$ makes all allocations stable. At the other extreme, the next two propositions illustrate the unconstrained power of an agenda setter when the alternative allocation protocol is unrestricted. Even with the incremental addition of small $\beta > 0$ (Proposition 3), the agenda setter retains the unconstrained power, as in the $\alpha = 1$ case (Proposition 2, from Thakur (2020)), to reach any alternative allocation from any initial allocation. This is akin to McKelvey’s Chaos Theorem (McKelvey, 1976): for $N \geq 5$ there exists a majority-approved sequence of reassignments from any allocation to any other allocation.

Proposition 2. *(Thakur, 2020) If $\alpha = 1$ and the alternative allocation protocol is unrestricted then for $N \geq 5$ there exists a majority-approved path of restructuring (sequence of reassignments) from any allocation to any other allocation.*

Proposition 3. *If $\alpha + \beta = 1$ and the alternative allocation protocol is unrestricted then for any $N \geq 5$ there exists a $\beta^* > 0$ such that for all $\beta < \beta^*$ there is a majority-approved path of restructuring (sequence of reassignments) from any allocation to any other allocation.*

Proof: Thakur (2020)’s proof (reproduced in Appendix B for the reader’s convenience) that for $N \geq 5$ there exists a majority-approved sequence of reassignments from any allocation to any other allocation implies that for each allocation on this path of restructuring, a majority is strictly better off. Hence, by continuity, there exists $\beta = \epsilon$ small enough such that this majority still approves each of the alternative allocations at every step on the path of restructuring.

Proposition 3 shows that Theorem 1’s conclusion — endogenous restructuring converging to and absorbed by the first best allocation — requires more than β being strictly positive; it must also assume that the alternative allocation protocol is restricted. Thus in this context, a procedural discipline (allowing only *constrained* endogenous organizational restructuring) complements a certain type of preference (agents valuing *organizational output*) in restraining the agenda setter’s power and ensuring that the process winds up at the maximally efficient matching of people to positions. As we emphasized in Section 2, Assumption A is maintained only for ease of exposition. However, analogous results about the path of restructuring becoming increasingly stable and converging to the first best hold under other alternative allocation protocols and other dynamic processes as well. For example, one could assume that transitions from one allocation to another are governed by a contest function, e.g., transition probabilities equal the fraction of voters preferring the alternative allocation to the status quo. Transitions to new allocations are then probabilistic, not deterministic. Although in this setting the first best is absorbing only for β sufficiently close to one (suboptimal allocations are never absorbing), the limiting distribution puts more probability on allocations that are more stable. The contest functions can represent the endogenous pressure for change based on the possibly informal power of different organizational coalitions. This approach does not presume that the organization literally votes (via some formal procedure) on agent-seat assignments; it is thus a relaxation of the voting idealization that we have used in the present paper.

What happens if agents have agenda setting power? Although Propositions 2 and 3 establish McKelvey (1976)-like chaos of majority voting, adding some structure highlights the robustness of our results about organizational efficiency in the limit, given $\alpha + \beta = 1$ and $\beta > 0$. Instead of Assumption A, for Theorem 2 below, we use *Assumption A***: i) for any status quo allocation in any period, the recognition probability that any agent i gets to propose the alternative proposal is strictly positive and stationary but can depend on the status quo allocation, and ii) from every status quo allocation, a recognized agent proposes an alternative allocation with positive probability if and only if it makes him strictly better off; if no such alternative allocation exists then

he proposes the status quo with probability 1. These assumptions produce a finite state Markov process with time-homogeneous transition probabilities.

Theorem 2. *If Assumption A^{**} holds and $\alpha + \beta = 1$ then myopic voting with majority rule implies the following: i) For every $\beta \in [0, 1]$, the Markov process has a unique stationary distribution and it converges to that limiting distribution from every initial probability vector over the set of allocations; ii) there exists $\underline{\beta} > 0$ such that for all $\beta < \underline{\beta}$, all allocations occur in the limiting distribution with positive probability; iii) there exists $\bar{\beta} < 1$ such that for all $\beta > \bar{\beta}$, the limiting distribution is degenerate, with all mass concentrated on the first best allocation; iv) the first best is the only allocation to occur with positive probability in the limiting distribution for every $\beta \in [0, 1]$.*

Proof: We have a finite Markov chain with stationary transition probabilities and Assumption A^{**} ii) guarantees aperiodicity. In Theorem 5i) we prove that the first best is uncovered for all $\beta \in [0, 1]$, hence it is recurrent for all $\beta \in [0, 1]$, as it is reachable from any other allocation under majority rule and myopic voting in one or two steps, which proves iv). This also implies that an allocation is either a transient state (to which there does not exist a majority-approved path of restructuring from the first best) or a recurrent state (to which there exists a majority-approved path of restructuring from the first best) that belongs to the single closed recurrent class, which proves i). For β small enough, Proposition 3 implies the Markov chain is irreducible as all states communicate (there exists a majority-approved path of restructuring from any allocation to any other allocation), hence every allocation is recurrent and occurs with positive probability in the limiting distribution, proving ii). Part iii) holds as the first best is reachable from any other allocation and, because it is the Condorcet winner for β sufficiently close to 1, it is therefore an absorbing state.

Theorem 2 allows for relatively unconstrained strategic agenda setting (e.g., just promoting himself or promoting nearly everybody) and agents with heterogeneous recognition probabilities who hence vary significantly in agenda-setting power.¹² Moreover, Assumption A^{**i}) also allows the recognition probabilities to depend on who is assigned to what position (e.g., executives or capable managers may have more say in proposing reshuffling). Despite this generality, the theorem highlights the privileged position of the first best allocation in the limiting distribution, thereby establishing the robustness of Theorem 1. Moreover, simulations show that the limiting distribution of this Markov process places the most weight on the first best allocation even for relatively small β . More generally, because allocations which produce more output have smaller win sets, they tend to be more stable. Although the unique first best allocation is not necessarily an absorbing state for very small β , it is the most stable allocation with the highest weight in the limiting distribution.

3.3. Informal Preference Aggregation: From Voting to Decisive Coalitions.

Theorem 1 obtained sharp results, but it is based on a stylized representation of organizational preference-aggregation where *every agent has equal power*. Moreover, it may appear that thus far, by modeling such voting *rules*, our paper only considers **formal procedures** within organizations for preference aggregation. In this section we show that one can analyze how conflicts over status and organizational performance get resolved with a more flexible notion that captures differential power and influence across agents. We do this via a representation of sets of decisive coalitions that easily lends itself to interpretations of **informal processes** of influence, persuasion, and negotiation. The results of this section show that the central thrust of Theorem 1 — even a slight weight on collective performance induces a restructuring path that converges to and is absorbed by the first best allocation — continues to hold under this more flexible interpretation.

¹²Both Assumptions A^{**i}) and A^{**ii}) can be relaxed still further without sacrificing the analytically useful framework of a finite Markov chain with stationary transition probabilities.

Given any two allocations A and A' , a coalition of agents $D \subseteq (1, \dots, N)$ is **decisive** if $U_i(A') > U_i(A) \forall i \in D$ implies that the collective choice is for A' over A .¹³ For any preference-aggregation process f that is a mapping from the set of agents' preferences to the set of reflexive and complete binary relations on the set of matchings, let $\mathbb{D}(f)$ be the set of **decisive coalitions** associated with f . The preference-aggregation process is thus defined by the (nonempty) set of decisive coalitions, i.e., the set of coalitions such that a unanimous agreement amongst all coalition members becomes the collective choice.

This generalized formulation includes many different preference-aggregation processes: dictatorship, oligarchy, majority rule, any weighted voting rule, and so forth. Within this large class a preference-aggregation process f is called collegial if and only if $K(D(f)) \equiv \bigcap_{D \in \mathbb{D}} D \neq \emptyset$. That is, an aggregation process is collegial if there exist agents, called a collegium $K(D(f))$, who appear in every decisive coalition. Thus, dictatorship and oligarchy are collegial. Any agent who has veto power would belong to the collegium. (The converse does not hold however: being in the collegium does not imply veto power.)

Clearly, a large set of preference-aggregation processes are **non-collegial**: no agent belongs to every decisive coalition. In such settings everyone loses sometimes. Agents may, however, differ enormously in their power and influence.

Despite this generality, the next result shows that provided that the alternative allocation protocol is restricted to incremental upside changes — $N_p = 1$ — Theorem 1's efficiency conclusions are retained.

Theorem 3. *If $\alpha + \beta = 1$, $\alpha, \beta > 0$, and no agent belongs to every decisive coalition then i) no suboptimal allocation is stable, and ii) the first best is stable and is dynamically reached with probability 1 if and only if the alternative allocation protocol is restricted to those in which $N_p = 1$.*

Proof: i) When no agent i belongs to every decisive coalition, any output-improving swap that involves demoting agent i cannot be blocked

¹³The set of decisive coalitions $\mathbb{D}$ is monotonic ($D \in \mathbb{D}$ and $\mathbb{I} \supseteq D' \supset D$, implies $D' \in \mathbb{D}(f)$) and proper ($D \in \mathbb{D}$ implies $\mathbb{I} \setminus D \notin \mathbb{D}$). See Austen-Smith and Banks (1999) for more discussion on decisive coalitions.

as all agents $\mathbb{I} \setminus i$ support it, and hence no suboptimal allocation is stable.

ii) $(\implies)$: That the first best is not stable given $N_p > 1$ is proven by an example. Without loss of generality consider the set of decisive coalitions $\{\{i, j\}, \{i, k\}, \{j, l\}\}$ for some agents i, j, k, l where i and j are not the agents indexed 1 and 2. This allows the agents i and j to defeat the first best for small enough β such that they value their status gain more than the output loss. $(\impliedby)$: Since alternative proposal are restricted to $N_p = 1$, $N_d + N_u = N - 1$. a) *first best is stable*: Since no agent belongs to every decisive coalition, the meritocratic first best allocation is stable because any promotion reduces output, thereby making all N_u uninvolved agents and all N_d demoted agents worse off. No other allocation is stable because every other set of person-position matches has at least one meritocratically twisted pair (such that $i < j$ despite $A(i) > A(j)$) and swapping such pairs makes $N - 2$ uninvolved agents and i better off. Since no agent belongs to every decisive coalition, agent j cannot unilaterally block the swap. b) *first best dynamically reached with probability 1*: Any promotion that reduces output would make all N_u uninvolved agents and all N_d demoted agents worse off, inducing them to oppose the proposed promotion; if output remained constant then all N_u uninvolved parties would be indifferent and all N_d demoted agents of the swapped agents would oppose. Hence, since no single agent belongs to every decisive coalition, the only proposals that pass are output-improving changes. This, plus the fact that the set of possible allocations is finite, together imply that the process must reach the first best with probability 1.

How much N_p must be constrained depends on the distribution of intra-organizational power. If power is equally shared (e.g., majority rule settles conflicts) then the constraint is just that a proposal may not promote a majority of the organization's personnel. If power is unequally distributed, as it can be in some non-collegial contexts, then N_p must be more sharply restrained to ensure that the first best allocation is stable. In all cases the underlying danger is that a group of sufficiently powerful agents will sacrifice organizational performance in order to improve their own status.

In fact, our result is robust to even more complex preference-aggregation processes in which power is based on which agent is assigned to which seat, i.e., allowing decisive coalitions to depend on the matching. For example, an agent’s influence over personnel reorganization can depend on whether he/she is a worker, a manager, or the CEO.

Theorem 4. *If $\alpha + \beta = 1$, $\beta > 0$, and no incorrectly matched (relative to his first best assignment) agent-seat pair belongs to every relevant decisive coalition¹⁴, then i) no suboptimal allocation is stable, and ii) the first best is stable and is dynamically reached with probability 1 if and only if the alternative allocation protocol is restricted to those in which $N_p = 1$.*

Proof: in Online Appendix A.

Theorems 3 and 4 underscore that Theorem 1’s conclusion — the path of restructuring gets absorbed into the first best allocation — does not require that agents have equal power or that conflicts are resolved by formal rules. (Thus majority rule, a formal procedure that encodes equal power, is not necessary.) Theorem 1’s conclusion is robust to alternative kinds of influence that arise organically, such as informal debate, deliberation, and position-based politicking and manipulation.

The intuition we have developed with non-collegial decisive coalitions carries over to other preference-aggregation processes, for example, weighted voting (where weights can even depend on agent-seat pairs, in line with Theorem 4). The analog of no agent (agent-seat pair) belongs to every (relevant) decisive coalition under weighted voting terms is that no single agent (agent-seat pair) has weight so large such that his/her consent is *always* needed for approving any proposed change.

¹⁴There are two technical notes: a) Relative to the status quo allocation, infeasible decisive coalitions are dropped when we consider the set of “relevant decisive coalitions”: i.e., when agent i is currently assigned to s , the set of relevant decisive coalitions is assumed to ignore decisive coalitions involving agent-seat pairs where i is assigned to any $s' \neq s$ and those where any j is assigned to s . b) From each possible allocation, the set of relevant decisive coalitions should be nonempty.

4. ORGANIZATIONAL IMPEDIMENTS TO REACHING EFFICIENCY.

4.1. Consequences of Friend Networks & Connections.

Friends, relatives, professional acquaintances, and other personal relationships all lead to complex networks of informal ties within an organization that can affect career opportunities (Podolny and Baron, 1997). These social networks can lead to alliances and payoff externalities that can affect how an organization evolves and how certain proposed changes are facilitated or blocked. For example, belonging to the same political faction (Francois et al., 2016), social ties based on hometown or college connections (Fisman et al., 2020), and caste-based similarities (Iyer and Mani, 2012) can affect promotions and assignments in a bureaucracy. This section speaks to the large literature on social capital theory and its interplay with incentive pay (see Oyer and Lazear (2007) and Lazear and Shaw (2007) for surveys of this literature) and issues of glass-ceilings and old boys' networks based on discrimination and segregation on gender, race, ethnicity, and so forth (see Smith (2002) for survey on gender and racial/ethnic minorities in the workplace).

For simplicity we call all of these social ties, “friends.” For tractability, we assume that each agent i assigns the same utility weight to all his friends α and $u_i = u$ for all i . We do not require that ties be symmetric: i can consider j his friend but j might not reciprocate. Asymmetric ties could arise for several reasons: e.g., although both i and j could have friendly feelings toward each other, j likes to insulate the workplace from personal ties whereas i is comfortable being friends with co-workers.

Intuitively when $\alpha + \gamma = 1$, with $\alpha, \gamma > 0$, friendships and other informal connections can help individuals climb an organization's status hierarchy, thereby impairing organizational efficiency via the promotion of unqualified but well-networked individuals. Luthans et al. (1988) found that the most successful managers spend 70 percent more time networking and 10 percent more time routinely communicating than do their counterparts. In this sense, a ‘useful’ friend is one who is not friends with many others: such agents are less likely to be friends with both the promotee(s) and the demotee(s) and thus avoid facing a conflict of friendship that may induce them to abstain.

If agents also care about organizational output ($\alpha + \beta + \gamma = 1$, with $\alpha, \beta, \gamma > 0$) then distortions arising from friendship ties can be alleviated. This helps to explain the productivity gains observed from individual or group-based incentive pay linked to output. For example, managerial incentive pay can lead to productivity improvements at the expense of the manager's friends who receive worse assignments under the incentive pay plan (Bandiera et al., 2008).

Crucially, however, stabilizing efficiency requires that the alternative allocation protocol must be further restricted.

Remark 4. *Assume that $u_i = u$ for all i , $\alpha + \beta + \gamma = 1$, with $\alpha, \gamma > 0$, $\beta \geq 0$ and the alternative allocation protocol is restricted to those in which $N_p < \frac{N+1}{2}$. A single asymmetric friendship of any agent i considering any other agent j to be his friend (except $j = 1$ or $i = 1$ and $j = 2$), can make the first best unstable for $\beta \geq 0$ small enough.*

Even the sparsest possible friendship network — just one directed link — renders the alternative allocation protocol used thus far ($N_p < \frac{N+1}{2}$) incapable of ensuring efficiency. This is because promoting a minority-sized coalition is also supported by the friend whose position is unchanged. Notice that the agent being friends of the highest quality agent ($j = 1$) does not undermine stability of the first best as there is no further promotion possible for such an agent.

The underlying coalitional logic of un-meritocratic promotees and their friends forming majority coalitions generalizes to adding more directed friendship links in support of different promotees or multiple friends supporting the same promotee. For example, agents i and j considering k to be their friend or agents i and j considering k and l to be their friends respectively can prevent alternative allocation protocols restricted to $N_p < \frac{N-1}{2}$ from stabilizing the first best. Continuing down this road leads to the conclusion that, given $\beta \geq 0$ small enough, the first best is not stable even when new allocation proposals are maximally restricted (i.e., to swaps) if any one agent $i \neq 1$ has too many (more than a majority) friends willing to support their un-meritocratic promotion despite the loss of some organizational output.

Of course, certain kinds of dense friendship networks do not overturn Theorem 1. For example, suppose everyone in the organization is friends of everyone else. This implies that the total utility anyone gets from the status of their friends is a constant, making all allocations equivalent on friendship grounds. Effectively, then, we are back to the $\alpha + \beta = 1$ setting and Theorem 1 still holds. Thus, the relation between the density of friendship networks and stability of the first best is not monotonic. Instead, Remark 4 simply illustrates how such interpersonal ties can help mobilize large winning coalitions to undermine the stability of the first best, even with the most stringent alternative allocation protocol restrictions.

4.1.1. *Cliques.*

In practice, friendship networks can often be *clique-based*: N agents are partitioned into a set of disjoint cliques such that every member of a clique is friends with everyone else in that clique—e.g., tight friendship circles or high-solidarity groups based on race, gender, ethnicity, or the like.

Intuitively, if a clique is ‘too big’ then it can get its way at the expense of everyone else. Thus, for sufficiently low β , a majority-sized clique will ensure that regardless of competence all its members get higher status positions than anyone outside their clique. (Members in the majority-sized clique will be ordered by competence, given $\beta > 0$.) Restrictions to $N_p = 1$ or swaps does not work in such a lopsided setting. Such networks can support tiered organizations: e.g., a male-dominated organization in which men are the majority and all occupy higher status positions than any woman regardless of quality (Matsa and Miller, 2011). This speaks to segregation, discrimination, sorting, and glass-ceiling effects within organizations based on informal groups and networks.

If there is no majority-sized clique then it turns out that requiring $N_p = 1$ is not only sufficient to make the first best stable, it is also necessary.

Proposition 4. *Assume that $u_i = u$ for all i , $\alpha + \beta + \gamma = 1$, with $\alpha, \beta, \gamma > 0$, and the friendship networks are clique-based. Then the first best allocation is stable for all clique-based networks that lack a majority-sized clique if and only if the agenda protocol is restricted to those in which $N_p = 1$.*

Proof: in Online Appendix A.

Thus cliques act like voting blocs and cleverly designed proposals can then construct a majority out of a set of minority voting blocs. Consequently, only the combination of restricting the alternative allocation protocol to $N_p = 1$ and constraining the size of the clique guarantee that the first best allocation is stable.

Overall, if friendship networks in an organization are dense and overlapping, then person-to-position assignments can often remain far from meritocratic and output can be far from optimal. This is consistent with the ample literature documenting inefficiencies in family firms.¹⁵ However, our analysis also highlights how an organization’s network of informal connections interacts with the alternative allocation protocol in ensuring stability against further reorganizations.

4.2. Consequences of Teams & Departments.

A nearly universal feature of organizations of even modest size is a team or departmental structure. It is therefore natural for agents to care about who else is in their team or department. Of course, if the team or department is given complete autonomy to organize its own personnel then a smaller version of our problem, restricted to this part of the organization, has been created. For example, discretion for personnel management is often given to franchises, subsidiaries, or decentralized companies of a large transnational corporation. In this section we model teams or departments that do not enjoy such complete autonomy in personnel allocation. The set of seats is partitioned into teams, hence each agent belongs to exactly one team depending on his assigned seat. The ζ term adds a pay-off to agent i based on the quality of the persons occupying the seats that constitute i ’s team. Having more capable individuals assigned to one’s team can yield numerous benefits: improving division of labor and efficiency, increasing group expertise both directly and indirectly (via informational spillovers), and improving work satisfaction due to the increased

¹⁵Many papers find a negative relationship between dynastic family CEOs and firm outcomes and productivity (Claessens et al., 2002; Perez-Gonzalez, 2006; Villalonga and Amit, 2006; Bertrand and Schoar, 2006; Bertrand et al., 2008; Cai et al., 2013; Bandiera et al., 2012; Lemos and Scur, 2018).

competence of colleagues. Teams are sometimes given group incentive pay, making the connection to caring about the quality of team members even more direct (Boning et al., 2007).

Remark 5. *Assume $\alpha + \beta + \zeta = 1$. A necessary but not sufficient condition for the first best allocation to be stable for all $\alpha, \zeta > 0$, $\beta \geq 0$ and all team structures is that the alternative allocation protocol is restricted to swaps.*

Remark 5 highlights that the presence of teams imposes even more demanding restrictions on alternative allocation protocols than does the presence of friendship networks. Given $N_p = 1$, departing from the first best can generate multiple demotions of more qualified agents that can benefit the teams they join (due to $\zeta > 0$, for β small enough). Moreover, these teams can collectively form a majority.

If the alternative allocation protocol is restricted to swaps and agents care about both status and team quality (i.e., $\alpha + \zeta = 1$, $\alpha > 0$, $\zeta > 0$), the logic of uninvolved agents extends to uninvolved *teams*. Given these parameter values, teams whose agents would be unaffected by the swap remain indifferent and abstain. Hence, if no team has at least a minimal majority-1 number of seats, all allocations are stable. If, however, a minimal majority+1 sized team¹⁶ exists then, because it votes to add the highest quality agents to itself, it may skew meritocratic assortment and produce inefficient organizational output.

Proposition 5. *If $\alpha + \zeta + \beta = 1$, $\alpha, \zeta, \beta > 0$ and the alternative allocation protocol is restricted to swaps then the following hold:*

- i) If there exists a team with at least minimal majority+1 seats,¹⁶ then in stable allocations every team is internally ordered by quality. However, depending on the parameter weights there may be agents in the majority+1 team who are under-placed.*
- ii) If no team has at least minimal majority-1 seats,¹⁶ then with probability 1 the first best is reached via output-improving swaps.*

Proof: in Online Appendix A.

¹⁶If there exists a team with a minimal majority or minimal majority-1 number of seats then it is consequential how the agents i, j involved in the swap vote, which depends on the exact coefficient weights and the functions u_i , u_j , f_i , and f_j (and $\mathbb{Y}$ for Proposition 5).

When agents also care about organizational output, uninvolved teams again vote for all output-improving swaps and vote against those that decrease output. As a result, all within-team seat assignments are meritocratically ordered. Moreover, the first best is reached if there is no majority team. However, a majority team can potentially skew meritocratic assortment and lead to sub-optimal organizational output by voting to attract higher quality agents who would have been more productive in another team.

These results highlight the potential harm caused by very large teams or departments within organizations. By monopolizing the vote over organizational changes, such teams can mis-allocate talent and produce sub-optimal output. For example, the dominance of the Xerox copier division is held responsible for the failure to effectively develop and commercialize Xerox PARC division's innovations of graphical user interface and personal computer, the Alto (Smith and Alexander, 1988). Though we do not consider the endogenous creation or dissolution of teams and departments, our discussion in this section has important implications for restructuring teams and departments after corporate mergers and acquisitions or other such significant changes that necessitate internal re-organization.

4.3. Discussion on Incremental Reorganizations.

Our results in the two previous subsections highlight that richer organizational features like friendship networks and teams/departments hinder the stability of the first best. Friendship networks and teams act as voting blocs with vested interests, thereby reducing the number of uninvolved agents from a proposed reorganization. Hence, to restore meritocratic discipline the organization must restrict the alternative allocation protocol to $N_p = 1$ given friend networks; given team/department structures it must go further, to the maximal restriction of pairwise swaps.

Pairwise swaps (i takes j 's seat and vice versa) of course satisfy $N_p < \frac{N+1}{2}$ requirement of Theorem 1. Hence without richer organizational features like friends and teams, output-improving swaps pass with an overwhelming $N - 1$ to 1 vote as the number of uninvolved agents is now $N_u = N - 2$.

Corollary 1.2. *If $\alpha + \beta = 1$, $\beta > 0$, and the alternative allocation protocol is restricted to swaps then the following hold: i) the path of restructuring is absorbed into the first-best—the unique stable allocation—with probability 1, ii) output is non-decreasing over the path of restructuring, and iii) the distance from the meritocratic ideal is weakly decreasing over the path of restructuring.*

Corollary 1.2 will often serve as the basis for comparison to illustrate the effects of other voting rules (Section 5) and sophisticated, forward-looking voting behavior (Section 6).

5. OTHER COMMON VOTING RULES

In this section we analyze the organization's path of restructuring under other common voting rules, contrasting it to Corollary 1.2 where the alternative allocation protocol is restricted to swaps.

First, because generalizing from majority rule to super-majority rules requiring Q votes, $\frac{N+1}{2} < Q < N$, does not affect the underlying logic of the N_u uninvolved agents voting only for output-improving reorganizations, all of the fundamental results continue to hold. Extending to super-majority rule just relaxes the alternative allocation protocol restriction in Theorem 1 to $N_p < Q$.

Second, under a unanimity rule all allocations are stable even if $\alpha + \beta = 1$, provided only that the following mild condition holds: β is not so large as to induce an agent to vote against his own status ranking in order to increase organizational output. Hence, the powerful effects of $\beta = \epsilon$ that hold under super-majority Q -rules, disappear under unanimity. Moreover, even with friends (unless i is friends with all other $N - 1$ agents) and departments, a very high weight on β is required to get any movement. Thus, the unanimity rule magnifies the impact of the fundamental rivalry inherent in choices over positions and status.

Third, consider a dictatorship, where agent i is the dictator. For $\alpha = 1$, all allocations where i gets the highest-status position are stable. With $\alpha + \beta = 1$ and $\alpha, \beta > 0$, the only allocation that is stable involves making dictator i optimally well off and everyone else ordered meritocratically. We say “optimally well off” because the dictator might not want himself in the highest-status position if the organization's loss of output would exceed his

gain in status. That is, the dictator puts himself at the highest seat such that if he were to increase to higher status position, $-\beta * \Delta Y > \alpha * \Delta status$. Although it is not surprising that concern over output leads to meritocratic assignments for agents with no power, it is interesting that sometimes output concerns can discipline even the dictator himself. Regarding friends, only having the dictator as a friend matters. Such relationships can skew the meritocratic ordering and decrease organizational output if the dictator values friendship enough. Regarding departments, only being in the dictator's unit matters; this can skew meritocratic ordering (the dictator's department works similarly to a majority+1 department). Note that a dictator whose actions are controlled by only short-term utility might take a path of restructuring which improves output but makes the allocation farther in terms of swap distance from the meritocratic ideal.¹⁷

Fourth, suppose some agent i has veto power. The most interesting case is when β , though strictly positive, is not so high that i will sacrifice status in order to increase organizational output. In this case, if i is in the highest status seat then he behaves like a dictator (with $\alpha + \beta = 1$). If his seat is below the top then he will veto any change that decreases either his status or organizational output. So once again, stable allocations are meritocratic apart from i 's own position. Moreover, if i starts off under-placed, he will be moved up because all N_u uninvolved agents vote for output-improving swaps. But if he is ever over-placed, either initially or as a result of a previous reorganization, then he will never go back down to his organizationally ideal position (unless of course he himself cares a great deal about output). Again, having i as a friend and hence enjoying i 's support protects an agent against being moved down the status hierarchy. Finally, being in the department of a veto player

¹⁷Suppose we have $N=15$. All agents are correctly assigned their ideal seat except #1, #2, #7, and #15. Seat 1 is assigned quality #7, seat 2 is assigned quality #15, seat 7 is assigned quality #1, and seat 15 is assigned quality #2. The fastest way to get to the first best is to swap agents 1 and 7 and agents 2 and 15 (i.e., two swaps). But a dictator with a 1-period time horizon might (depending on the output function) switch agents 1 and 15 in seats 2 and 7. From this allocation, it takes a minimum of 3 swaps to get to the first best. This is relevant for short-run stock option compensation of CEOs which can incentivize acting against long-run growth via short-run manipulations.

ensures that one's unit can never be made worse off even if it is a minority department.

Fifth, consider an oligarchy (of two or more agents) defined as a set of agents G such that all $i \in G$ have a veto and are collectively decisive: an alternative A' is passed if and only if all $i \in G$ prefer $A' \succ_i A$. If $\alpha = 1$ then everything is stable when the alternative allocation protocol is restricted to swaps. Given $\alpha + \beta = 1$ and $\alpha, \beta > 0$, all meritocracy-improving swaps are approved unless they displace a member of the oligarchy (such proposals are vetoed). Oligarchs can move up only if a) they do not take another oligarch's seat and b) output goes up. Given certain starting allocations, oligarchs may be unable to reach the highest status seats. If, however, they hold those seats in the initial allocation then they cannot be displaced from them. All other seats are meritocratically ordered.

The big difference between oligarchy and dictatorship is that a dictator has no challengers. Someone that dominant will give himself as much status as is egoistically optimal (his optimal position depends on the β weight), but he will be ruthlessly meritocratic regarding everyone else in the organization. In contrast, oligarchs are rivalrous against each other: though they all vote to impose meritocratic discipline on everyone else, they do not want to endure such discipline themselves and they will veto meritocracy-enhancing proposals that would hurt them. Thus, they might block each other's movements given $\beta > 0$. Hence for β strictly positive but not so large that everyone unconditionally prefers meritocracy, under-placed oligarchs improve their own positions if output increases, all the way up to their ideal (unless it is occupied by another oligarch who will not move), and an over-placed oligarch vetoes being made worse off.¹⁸ The privileged circumstances of dictators and oligarchs — disciplining others but not necessarily holding themselves accountable — can be a reality amongst the executive levels of many organizations; for example, Longenecker and Gioia (1988) find that 40 percent of interviewed executives reported not receiving annual performance appraisals.

¹⁸Hence, if a shock to the organization's task environment changes the agents' quality ordering (per the discussion following Corollary 1.1), newly over-placed oligarchs can block adjustments that would move the organization to the new post-shock meritocratic ordering.

Next, consider the organization using a plurality rule. The logic of unininvolved agents supporting output-improving moves is still robust for $\beta > 0$; what changes is what can happen given $\beta = 0$. In that case a few friends (with $\alpha + \gamma = 1$) or even non-majority departments of different sizes (with $\alpha + \zeta = 1$) can lead to reorganizations, passed by a plurality, which are neither meritocracy- nor output-improving. With friends and $\alpha + \gamma = 1$, seats are assigned in order of status to agents based on how many friends they have. Regarding departments, if $\alpha + \zeta = 1$ then the quality of those assigned to a department's seats is based on the relative size of the department. Hence, while the unininvolved agents will, given $\beta > 0$, vote for output-improving moves, thus overwhelming plurality rule, friends and departments can pass the lower threshold (of being relatively more connected or wanted by a larger department) to make certain moves possible and certain allocations unstable even when $\beta = 0$.

Plurality rules enable group-based hierarchies. Under plurality rule, if $\alpha + \gamma = 1$, $\alpha, \gamma > 0$, allocation protocol is restricted to swaps, and friendship is symmetric (i.e., i is a friend of j implies j is a friend of i), then all allocations where seats are assigned in order of status to agents in order of the absolute number of friends they have are stable and every path of restructuring converges with probability 1 to one of these stable allocations. Moreover, if the set of friends forms an equivalence class (i.e., in any particular group of friends everyone is friends with each other) then there is a strict dominance hierarchy based on membership in one's friendship bloc and every path of restructuring converges to one of these strict dominance hierarchies with probability 1. Further, within a friendship bloc every allocation is stable.

Finally, in line with certain personnel decisions sometimes being “above a certain pay-grade,” we can also allow voting to be restricted to seats above a status threshold.¹⁹ Specifically, when voting is restricted to agents in seats above a pay-grade threshold $\underline{s}$, only those agents in seats $s \leq \underline{s}$ can vote or

¹⁹Convergence to the first best under $\alpha + \beta = 1$, $\alpha, \beta > 0$ and the alternative allocation protocol being restricted to swaps (Corollary 1.2) is robust to many other formulations of weighted voting, such as restricting voting to positions nearby the swapped positions (e.g., only the relevant managers and other employees working directly under an employee are often asked for performance appraisals used to consider promotions), or allowing for the possibility that many agents have no influence or say in certain personnel changes.

have a say. Moreover, swaps involving different seats i, j may have different status thresholds $\underline{s}(i, j)$: e.g., a decision to replace a CEO might be restricted to the Board of Directors, whereas replacing a mid-level manager may trigger the involvement only of the team working directly under that manager and his/her immediate superior.

Remark 6. *With voting restricted to agents in seats above a pay-grade thresholds such that $\min_{i,j} \underline{s}(i, j) > 2$ for swaps involving any two seats i and j , if $\alpha + \beta = 1$, $\beta > 0$, and the alternative allocation protocol is restricted to swaps then the first best is reached with probability 1.*

Namely, with $\underline{s} \geq 3$, any output-improving swaps involving seat(s) $s \leq \underline{s}$ is approved by $\underline{s} - 1$ to 1 vote, any output-improving swap of positions $s > \underline{s}$ is supported unanimously, and any output-decreasing swap is voted against by at least a $\underline{s} - 1$ to 1 vote. Else, there is too much concentration of voting power at the top.

6. TOWARDS MORE SOPHISTICATED VOTING

Thus far we have assumed myopic voting. However, the strong organizational drift towards higher output (provided agents do not put zero weight on output) does not hinge on that assumption. In this section we show that Corollary 1.2, where the alternative allocation protocol is restricted to swaps, is robust to a variety of more sophisticated strategic voting behaviors.

6.1. Forward-looking opportunity costs.

When agent i considers an output-improving swap involving agents j and k , how is i 's set of potential swaps next period changed? Agents may vote against a swap today if it limits the opportunities for their own improvement going forward. This is 1-period forward-looking behavior regarding the set of possible actions in the next period.

When considering swapping agents j and k in the current period with $\beta > 0$, the set of attainable swaps in next period is changed only for those agents whose quality lies between j and k and who occupy seats in the status hierarchy in between or inferior to j and k 's seats. Let B denote this 'between' set. Agents in set B lose out on the opportunity to swap with agent k who is the

lower quality agent occupying a higher status seat because after the current period's swap is executed, agents in B cannot swap with agent j who is of a higher quality than they are.

Such swap-optionalty based on forward-looking behavior disciplines the organization from correcting 'too quickly' by putting the very deserving high-quality agents immediately on top. This is because all the agents in between lose out on the chance to challenge the 'weak' over-placed agent in the high status seat. Convergence to the first best is nevertheless guaranteed as output-improving swaps between agents m and n with $n - m < \frac{N+1}{2}$ difference in quality are always passed by majority rule as there are too few agents in the corresponding set B to vote against. Hence, this type of forward-looking behavior can only slow down convergence to the first best (given $\beta > 0$ of course).

6.2. Combining short-run and long-run considerations.

Consider a reduced form approach to agents combining both short-run and long-run evaluations when voting with $\alpha + \beta = 1$ and $\alpha, \beta > 0$. Agent i votes for a proposed output-increasing change if and only if (1) the alternative allocation proposed for period $t+1$ is strictly better for i than is the status quo allocation in period t , and (2) the first best allocation is strictly better for i than is the status quo allocation in period t . We refer to such voting behavior as *short-/far-sighted voting*. Assume that passing an alternative allocation proposal requires majority support. Let A_s denote the subset of all allocations that are stable and let $A_s(\beta)$ refer to the stable subset of allocations given a particular value of β .

Proposition 6. *Assume short-/far-sighted voting. If $\alpha + \beta = 1$, $\alpha, \beta > 0$, and the alternative allocation protocol is restricted to those in which $N_p < \frac{N+1}{2}$, then the following hold:*

- i) There exists $\beta^* \in (0, 1)$ such that if $\beta > \beta^*$ then the only stable allocation is the first best and it is reached with probability 1, and if $\beta < \beta^*$ then there exist inefficient stable allocations.*
- ii) If $\beta_2 > \beta_1 > 0$ then $A_s(\beta_2) \subseteq A_s(\beta_1)$.*

Proof: see Appendix A.

The first best is always stable under short-/far-sighted voting when agents value organizational output (i.e., $\beta > 0$). It is uniquely stable and reached with probability 1 for sufficiently large β . If inefficient stable allocations exist, then the path of restructuring gets hung up at the first stable inefficient allocation it encounters or it is absorbed into the first best.

The path of restructuring gets hung up at inefficient allocations with positive probability when status is sufficiently important (α is sufficiently high). The only inefficient allocations that can be stable are those in which a majority of agents enjoy positions of higher status than they would be occupying in the first best arrangement. If β is sufficiently low then the payoff from the private gain of status outweighs the collective loss of organizational output, producing majority support for an inefficient allocation. All other inefficient allocations can, as we have seen, be destabilized by pairwise swaps (amongst other arrangements).

Inefficient stable allocations can be far from the first best in terms of both output and meritocratic distance. Agents being far-sighted thus has potentially unfortunate consequences for efficiency: far-sighted agents anticipate the organizational pull towards meritocracy and can strategically vote against reorganizations that may ultimately come at their own expense.

Importantly, whether the path of restructuring gets hung up is both path-dependent and sensitive to the initial status quo. Short-/far-sighted voting highlights that meritocracy is a discipline we want to impose on others but which we dislike when it hurts ourselves. Far-sighted voting can shrink the set of uninvolved agents N_u who realize that ultimately, they too have skin in the game and can suffer from the imposition of meritocratic discipline.

6.3. Using uncovered sets to capture sophisticated voting.

As we know, there is no Condorcet winner when β isn't sufficiently close to one: even the first best allocation can be beaten by proposed allocations that would improve the status of a great many agents. In such circumstances sophisticated voting may be quite complicated, both for the agents to compute and for modellers to analyze.

Miller (1980) argued that we could gain significant understanding about what outcomes would be generated by sophisticated voting in these complex

settings by restricting our attention to the set of strategically undominated (‘uncovered’ was his term) alternatives. Adapting these notions to our setting, an allocation A *covers* allocation B if (i) $A \succ B$ and (ii) if allocation $C \succ A$ then $C \succ B$. An allocation A is *uncovered* if for all allocations B , either $A \succ B$ or there exists allocation C such that $A \succ C$ and $C \succ B$. Hence, an uncovered allocation beats every other allocation either directly or indirectly in two steps — a strategically important property.

We assume majority rule is used to compare alternatives. Note that the majority rule relation is strict: either the alternative is instituted with a strict majority or, if there is no strict majority favoring the alternative, the status quo is retained.

Miller’s recommendation to focus on the uncovered set has been supported by its robust existence even in chaotic environments (e.g., the Hobbesian world of $\beta = 0$) and by its nice properties: among other things, it collapses to the Condorcet winner when one exists, it is a subset of the pareto and top cycle sets, and it is often confined to a comparatively small subset of the entire set of alternatives. Most importantly, a variety of competitive choice processes, including sophisticated voting under standard amendment procedure, cooperative voting with free coalition formation, and open agenda formation, to name a few, have been shown to produce outcomes in the uncovered set. (See Miller [2007] for more on the properties of and results related to the covering relation and the uncovered set.) Accordingly, we analyze the covering relation and the uncovered set to examine the implications of sophisticated voting in our matching setting.

Theorem 5. *If $\alpha + \beta = 1$ then i) the first best allocation is always in the uncovered set for $\beta \in [0, 1]$, ii) if $\beta = 0$ then all allocations are in the uncovered set, iii) there exists $\beta^* < 1$ such that for all $\beta > \beta^*$, allocation A covers B if and only if $\mathbb{Y}(A) > \mathbb{Y}(B)$ and the first best is the only allocation in the uncovered set, and iv) the worst output allocation covers no other allocation for $\beta \in [0, 1]$.*

Proof: in Online Appendix A.

Theorem 5 highlights the close relation between output-efficiency and sophisticated strategic voting. However, this relation, though strong, is not simple: in particular, one cannot strengthen the necessity component of part iii) to “for all β , if allocation A covers B then A ’s output must exceed B ’s.” That claim is false.²⁰ However, simulations suggest that such counterexamples are rare.²¹

7. CONCLUSIONS

In this paper we analyze how agents’ preferences, especially for personal status and organizational output, affect the endogenous assignments of agents to positions within an organization. We find that when agents care even slightly about organizational performance, the organization moves toward meritocratic placement and efficient output. In practice, owning shares of a company or having some intrinsic skin in the game, e.g., personnel decisions in a wartime military or political committee assignments during legislative battles between political parties, can induce agents to value organizational performance. This commonality of preference for improving organizational performance makes otherwise uninvolved parties, whose own assignments remain unchanged, support output-improving changes in allocations. The organization reaches meritocratic, output-maximizing allocations in the limit as the stable, absorbing state. We show that even when agents are relatively sophisticated in anticipating where the path of restructuring is headed and vote strategically, the meritocratic, efficient assignment retains its privileged position of stability.

Our model reveals the importance of the *procedures* governing the kinds of allocations that can be proposed as alternatives to the status quo and how

²⁰Consider the following counterexample: $N = 5$, agents’ qualities are parameterized by 5,4,3,2,1 where a higher quality is better. The status of the 5 seats is parameterized by 5,4,3,2,1. $\mathbb{Y}(A) = \sum_{i \in N} (\text{SeatStatus}_i)^2 * \text{Quality}_i$, $\beta = 0.7$, and all agents get utilities [50, 40, 30, 20, 10] from being assigned to those seats respectively. Notation: we write an allocation in terms of what status seat is given to each person, ordered by decreasing quality (e.g., 45321 means the highest quality agent is given the seat of status 4, the second-best agent gets the highest status seat 5, and so on). Here notice that allocation $B = 45321$, which produces output 216, is covered by allocations $A_1 = 54132, A_2 = 53412, A_3 = 54321, A_4 = 54312, A_5 = 54231, A_6 = 53421$ which produce outputs 214, 215, 225, 222, 220, and 218.

²¹We also conjecture that the size of the uncovered set is weakly decreasing in β . Simulations have revealed no counterexamples thus far.

these protocols interact with the agents' *preferences* over positions and agents' *power* to institute change. Lacking any common interest in organizational output, perfectly correlated preferences on private goods — what constitutes higher status positions — generates so much competition, disagreement, and rivalry that no allocation is stable if alternative allocations are unrestricted. We highlight the interplay between the voting rule and the restrictions on the scope and size of reorganization necessary for stability. Without such restrictions, large endogenous coalitions keep pushing for self-serving, anti-meritocratic promotions for themselves at the expense of organizational efficiency, thus generating a debilitating systemic instability. As such, we study the full range of restrictiveness on the alternative allocation protocol, from allowing any alternative allocation to the maximal restriction of only allowing pairwise swaps.

We also highlight other dynamic complexities and practical impediments to efficiency-enhancing paths of restructuring. Firstly, very large departments, dominant friendship cliques, and extremely well-connected individuals can enlist large voting blocs to oppose meritocratic changes that would improve organizational performance. These richer organizational features increase the number of agents having vested interests in proposed reorganizations, thereby reducing the number of uninvolved agents. We demonstrate how restoring efficiency then necessitates constraining the set of alternative allocations that can be proposed. Secondly, individuals with undue influence in deliberation and collective decision-making can further their self-interest by blocking efficiency-enhancing moves.

Making any team, department, or group within the organization too large can skew talent and generate sub-optimal output. And extremely well-networked individuals can use their connections to improve their position within an organization even though their climb to high status subverts meritocracy, thereby degrading organizational performance. Individuals can use their power and influence to insulate their own positions from being reassigned through reorganizations. Even so, all such influential individuals want meritocratic discipline imposed *on other people*. For example, dictators in our model, like Enlightenment monarchs, want everyone else to be meritocratically assigned. This

points up the collective action problem inherent in the situation: few people are content when meritocracy hurts their own careers, but we all want competent doctors. Recent criticisms of meritocracy (e.g., Markovits, 2019; Sandel, 2020) tend to overlook this strategic complexity.

Both the instability generated by large coalitions driven by status-seeking and the inertia produced by incremental efforts at restructuring blocked by status-protection are empirical possibilities, not just theoretical properties. However, the main thrust of our model, which we think plays a crucial role in practice, is the strong stochastic drift toward meritocracy produced by *uninvolved parties* benefiting from meritocratically ordered assignments in organizations.

REFERENCES

- [1] Abraham, David J. Irving, Robert W., Kavitha, Telikepalli, and Kurt Mehlhorn. "Popular matchings." *SIAM Journal on Computing*, 37(4) (2007): 1030-1045.
- [2] Akerlof, George A., and Rachel E. Kranton. "Identity and the Economics of Organizations." *Journal of Economic perspectives* 19.1 (2005): 9-32.
- [3] Austen-Smith, David, and Jeffrey S. Banks. *Positive political theory*. Vol. 2. University of Michigan Press, 1999.
- [4] Bandiera, Oriana, Iwan Barankay, and Imran Rasul. "Social capital in the workplace: Evidence on its formation and consequences." *Labour Economics* 15.4 (2008): 724-748.
- [5] Bandiera, O., Guiso, L., Prat, A. and Sadun, R. (2012), What do ceos do?, CEP Discussion Papers dp1145, Centre for Economic Performance, LSE.
- [6] Barbera, Salvador, and Matthew O. Jackson. "Choosing how to choose: Self-stable majority rules and constitutions." *The Quarterly Journal of Economics* 119.3 (2004): 1011-1048.
- [7] Bertrand, Marianne, et al. "Mixing family with business: A study of Thai business groups and the families behind them." *Journal of financial Economics* 88.3 (2008): 466-498.
- [8] Bertrand, Marianne, and Antoinette Schoar. "Managing with style: The effect of managers on firm policies." *The Quarterly Journal of Economics* 118.4 (2003): 1169-1208.
- [9] Boning, Brent, Casey Ichniowski, and Kathryn Shaw. "Opportunity counts: Teams and the effectiveness of production incentives." *Journal of Labor Economics* 25.4 (2007): 613-650.
- [10] Cai, Hongbin, et al. "Family ties and organizational design: Evidence from Chinese private firms." *Review of Economics and Statistics* 95.3 (2013): 850-867.

- [11] Claessens, Stijn, et al. "Disentangling the incentive and entrenchment effects of large shareholdings." *The Journal of Finance* 57.6 (2002): 2741-2771.
- [12] Core, John, and Wayne Guay. "The use of equity grants to manage optimal equity incentive levels." *Journal of Accounting and Economics* 28.2 (1999): 151-184.
- [13] Cowgill, Bo, and Rembrand Koning. "Matching Markets for Googlers. Harvard Business School Case 718-487, March 2018. (Revised August 2018.)
- [14] Cseh, Agnes. "Popular matchings." *Trends in Computational Social Choice* (2017): 105.
- [15] Demsetz, Harold, and Kenneth Lehn. "The structure of corporate ownership: Causes and consequences." *Journal of Political Economy* 93.6 (1985): 1155-1177.
- [16] Francois, Patrick, Francesco Trebbi, and Kairong Xiao. *Factions in nondemocracies: Theory and evidence from the Chinese Communist Party*. No. w22775. National Bureau of Economic Research, 2016.
- [17] Fisman, Raymond, et al. "Social Ties and the Selection of China's Political Elite." *American Economic Review* 110.6 (2020): 1752-81.
- [18] Hannes, Sharon. "Reverse Monitoring: On the Hidden Role of Employee Stock-Based Compensation." *Mich. L. Rev.* 105 (2006): 1421.
- [19] Himmelberg, Charles P., R. Glenn Hubbard, and Darius Palia. "Understanding the determinants of managerial ownership and the link between ownership and performance." *Journal of financial economics* 53.3 (1999): 353-384.
- [20] Iyer, Lakshmi, and Anandi Mani. "Traveling agents: political change and bureaucratic turnover in India." *Review of Economics and Statistics* 94.3 (2012): 723-739.
- [21] Pérez-González, Francisco. "Inherited control and firm performance." *American Economic Review* 96.5 (2006): 1559-1588.
- [22] Jensen, Michael C., and William H. Meckling. "Theory of the firm: Managerial behavior, agency costs, and ownership structure." *Economics social institutions*. Springer, Dordrecht, 1979. 163-231.
- [23] Koray, Semih. "Self-selective social choice functions verify Arrow and Gibbard-Satterthwaite theorems." *Econometrica* 68.4 (2000): 981-996.
- [24] Krehbiel, Keith. *Information and legislative organization*. University of Michigan Press, 1992.
- [25] Lambert, Richard A. "Executive effort and selection of risky projects." *The RAND Journal of Economics* (1986): 77-88.
- [26] Lazear, Edward P., and Kathryn L. Shaw. "Personnel economics: The economist's view of human resources." *Journal of economic perspectives* 21.4 (2007): 91-114.
- [27] Lazear, Edward P., and Paul Oyer. *Personnel economics*. No. w13480. National Bureau of economic research, 2007.
- [28] Lazear, Edward P., and Sherwin Rosen. "Rank-order tournaments as optimum labor contracts." *Journal of political Economy* 89.5 (1981): 841-864.

- [29] Lemos, Renata, and Daniela Scur. "All in the family? CEO choice and firm organization." (2018).
- [30] Longenecker, Clinton O., and Dennis A. Gioia. "Neglected at the top-executives talk about executive appraisal." *MIT Sloan Management Review* 29.2 (1988): 41.
- [31] Luthans, Fred, Richard M. Hodgetts, and Stuart A. Rosenkrantz. *Real managers*. Ballinger, 1988.
- [32] Maggi, Giovanni, and Massimo Morelli. "Self-enforcing voting in international organizations." *American Economic Review* 96.4 (2006): 1137-1158.
- [33] Manlove, David. *Popular Matchings*, Chapter 7 in *Algorithmics of matching under preferences*. Vol. 2. World Scientific, 2013.
- [34] Markovits, Daniel. *The meritocracy trap: How America's foundational myth feeds inequality, dismantles the middle class, and devours the elite*. Penguin Books, 2020.
- [35] Matsa, David A., and Amalia R. Miller. "Chipping away at the glass ceiling: Gender spillovers in corporate leadership." *American Economic Review* 101.3 (2011): 635-39.
- [36] Mayhew, David R. *Congress: The electoral connection*. Yale university press, 1974.
- [37] McKelvey, Richard D. "Intransitivities in multidimensional voting models and some implications for agenda control." (1976).
- [38] Messner, Matthias, and Mattias K. Polborn. "Voting on majority rules." *The Review of Economic Studies* 71.1 (2004): 115-132.
- [39] Miller, Nicholas R. "A new solution set for tournaments and majority voting: Further graph-theoretical approaches to the theory of voting." *American Journal of Political Science* (1980): 68-96.
- [40] Pérez-González, Francisco. "Inherited control and firm performance." *American Economic Review* 96.5 (2006): 1559-1588.
- [41] Podolny, Joel M., and James N. Baron. "Resources and relationships: Social networks and mobility in the workplace." *American Sociological Review* (1997): 673-693.
- [42] Rajgopal, Shivaram, and Terry Shevlin. "Empirical evidence on the relation between stock option compensation and risk taking." *Journal of Accounting and Economics* 33.2 (2002): 145-171.
- [43] Sandel, Michael J. *The tyranny of merit: What's become of the common good?*. Penguin UK, 2020.
- [44] Smith, Ryan A. "Race, gender, and authority in the workplace: Theory and research." *Annual Review of Sociology* 28.1 (2002): 509-542.
- [45] Smith, Douglas K., and Robert C. Alexander. "Fumbling the Future: How Xerox Invented. Then Ignored, the First Personal Computer," William Morrow and Co, New York, NY (1988).
- [46] Sonmez, Tayfun, and Tobias B. Switzer. "Matching With (Branch-of-Choice) Contracts at the United States Military Academy." *Econometrica* 81.2 (2013): 451-488.

- [47] Thakur, Ashutosh. “Combining social choice and matching theory to understand institutional stability.” (2020).
- [48] Thakur, Ashutosh. “Matching Politicians to Committees.” (2017).
- [49] Thakur, Ashutosh. “Matching in the Civil Service: A Market Design Approach to Public Administration and Development” (2019)
- [50] Villalonga, B. and Amit, R. (2006), ‘How do family ownership, control and management affect firm value?’, *Journal of Financial Economics* 80, 385–417
- [51] Weitzman, Martin. “The Share Economy: Conquering Stagflation.” Cambridge, MA: Harvard University Press (1984)

ONLINE APPENDIX A. PROOFS

Proof of Corollary 1.1: If output would decrease under the new proposal then all N_d demoted agents and N_u uninvolved agents (hence, a majority) would vote against the proposed change. If output would remain constant then all N_d uninvolved parties would be indifferent and abstain while N_d demoted agents vote against. Hence the only proposals that are approved are output-improving changes where the uninvolved parties support the new allocation.

Proof of Proposition 1: $(\implies)$: Suppose the alternative allocation protocol admits challenges with $N_u < \frac{N+1}{2}$. Then consider the first best allocation to be the status quo in any period. Construct agents' preferences over seats such that each agent i prefers seat $i - 1$ to seat i for $i = 2, 3, \dots, \frac{N+1}{2}$ and agent 1 prefers seat $\frac{N+1}{2}$ to seat 1. Then for $\beta > 0$ small enough, all agents 1 through $\frac{N+1}{2}$ vote for the alternative allocation that moves agents $2, 3, \dots, \frac{N+1}{2}$ up one seat and agent 1 to seat $\frac{N+1}{2}$ over the first best. Although the uninvolved would vote against this change, they are a (bare) minority. $(\impliedby)$: For any allocation that produces less than maximal output, the uninvolved agents, who form a majority, vote in favor of (against) any output-improving (-decreasing) swap. Hence, since there are finitely many allocations, the allocation that maximizes output is reached with probability 1. This output-maximizing allocation is stable because every proposed alternative arrangement decreases output, inducing the uninvolved agents to block any such proposal.

Proof of Remark 3: We will prove this result by constructing a counterexample to the conjecture that the distance from the meritocratic ideal is weakly decreasing in every period of any path of restructuring if the assumptions of Theorem 1 hold. In the counterexample there are seven agents. Consider their seat-assignments: let the status quo be $J = (5, 2, 1, 3, 6, 4, 7)$, i.e., agents 5, 2, ... assigned to seats 1, 2, ..., and alternative allocation be $K = (2, 4, 5, 1, 3, 7, 6)$. Hence going from J to K , agents 2, 4 and 7 improve status while agents 5, 1, 3, and 6 lose status. Let $\mathbb{Y} = \sum_i (N + 1 - i) * (N + 1 - A(i))$. J generates output 127 while K generates output 128, hence there exists $\beta > 0$ such that K is majority approved. However, the distance from the meritocratic ideal is 10 for J , while it is 12 for K .

Proof of Theorem 4: Same as Theorem 3, replacing “agent” with “incorrectly matched agent-seat pair.” It is fine for the agent who is correctly matched to the seat he gets in the first best to be in every decisive coalition since with $\beta > 0$, he always votes to approve swaps that increase output and always votes against those that decrease output. Technical note a) guarantees that starting from any particular status quo allocation, the set of decisive coalitions does not include coalitions which include infeasible agent-seat pairs (i.e., only consider “relevant decisive coalitions”). Technical note b) guarantees that there is a nonempty set of relevant decisive coalitions at each status quo allocation, else the system could not leave any such a status quo.

Proof of Proposition 4: $(\implies)$: That the first best is not stable given $N_p > 1$ is proven by an example. Consider a proposal promoting agent i from clique 1 and agent j from clique 2 and suppose the two minority-sized cliques together form a majority. Then for β sufficiently low everyone from clique 1, being friends of i , will vote for the proposal; similarly, everyone from clique 2, as friends of j , will vote for it. $(\impliedby)$: Since the alternative proposals are restricted to $N_p = 1$, $\beta > 0$, and all alternative allocations produce less output than the first best, the only coalition that can support an alternative allocation can involve the promotee’s clique. But no clique is majority-sized.

Proof of Proposition 5: We will use the shorthand “majority+1 team” to stand for “team with at least minimal majority+1 seats.”

- Proof of i): First, consider a within-team swap. A proposal will pass if and only if it increases output; the $N - 2$ uninvolved agents vote for such changes and against proposals that would reduce output. Second, consider a swap between two minority teams. If that would decrease organizational output, the uninvolved majority+1 team will vote against; if output would rise, the uninvolved majority+1 team will vote for. Hence only swaps that increase output are passed. Lastly, consider a swap between the majority+1 team and some minority team. All uninvolved teams will only support output-improving swaps. The overall voting behavior depends on the relative weights ζ and β , the output production function $\mathbb{Y}$, and functions u_i and f_i . There are several cases

to consider:

- a) The majority+1 team will always vote for a swap that improves the quality of its own team and increases output. Hence such swaps always pass.
- b) The majority+1 team will never vote for a swap that impairs the quality of its team and decreases output. Hence such swaps never pass.
- c) A swap that decreases majority+1 team quality and increases output passes only if the majority+1 team values output enough.
- d) All minority teams, both involved and uninvolved, will never vote for a swap that improves the quality of the majority+1 team and decreases output, but the majority+1 team might vote for such a swap if it values team quality enough. In this case the meritocratic allocation might not be an absorbing state because the majority+1 team could have an incentive to skew the distribution to attract better talent in its own team.
- Proof of ii): First, consider a within-team swap. Only those that improve output will pass; these are supported by the $N - 2$ uninvolved agents. Second, consider a swap between two minority teams. Uninvolved teams will vote for the swap if and only if output increases. Moreover, if output were to decrease, then the involved team which got the lower quality agent from the swap would also vote against. Thus, only output-improving swaps are passed. Given that there are only finitely many possible matchings, the process converges with probability 1 to the first best allocation.

Proof of Proposition 6: Because we will use part ii) to prove i), we prove ii) first. Suppose that for β_2 allocation A , though suboptimal, is stable: a majority of agents prefer it to the first best allocation. This means that for those agents, the payoff from increased status more than compensates for the shortfall in organizational output. Next consider $\beta_1 < \beta_2$. Now even more weight is put on status, so again, the shortfall in organizational output must be more than compensated for by the private payoff in status. Therefore, if

an allocation is stable given β_2 then it must also be stable given β_1 , hence $A_s(\beta_2) \subseteq A_s(\beta_1)$.

Now turn to part i). Set $\beta = 1$. Then all agents strictly prefer the first best to any other allocation. Since this is strict it continues to hold for β sufficiently close to 1. Hence within this parametric region the far-sighted criterion is always satisfied, every output-improving change will be passed, and the path of restructuring converges to the first best with probability 1.

Now set $\beta = 0$. If the status quo allocation is $(2, \dots, N, 1)$, i.e., the most competent agent is given the bottom seat and everybody else moves up one seat, then $N - 1$ agents strictly prefer this to the first best allocation. Hence they must continue to prefer it for β sufficiently close to 0. To see this, note that if an alternative allocation produces less or equal output than the status quo then only the $N_p < \frac{N+1}{2}$ agents who would be promoted will vote for it; hence it fails. If, on the other hand, the alternative produces more output than the status quo then the far-sighted criterion kicks in: all the agents who prefer the status quo to the first best vote against the proposed allocation, and since they are a majority the new alternative is rejected.

Thus, if β is sufficiently big then the process converges to the first best with probability 1 while if β is sufficiently small then an inefficient stable allocation exists. These two facts, combined with ii), imply the existence of β^* .

Proof of Theorem 5:

- Proof of ii): When $\beta = 0$, all possible assignments lead to an equivalence with respect to relabeling agents, as only status matters (i.e., $\alpha = 1$) and the environment is completely rivalrous given the common status hierarchy. So relabeling makes any two alternatives an isomorphism. Hence, either all assignments are covered or they are all uncovered. However, the former is impossible since in every finite set of alternatives, at least one alternative is uncovered (Miller, 1980). Hence all allocations are uncovered.
- Proof of i): For $\beta = 1$, the first best allocation is a Condorcet winner as it is unanimously preferred over any other allocation. It therefore belongs to the uncovered set. For $\beta = 0$, the proof of ii) above implies the first best is uncovered in that case as well.

Consider, therefore, the remaining case of $\beta \in (0, 1)$. Label the first best allocation X . To show that X is uncovered, we will show that for any allocation Y , either $X \succ Y$ or there exists an allocation Z such that $X \succ Z$ and $Z \succ Y$.

For any alternative allocation Y , relative to X , there are agents whose seats remain unchanged (“unchanged”), agents whose seats are given worse status seats in Y than in X (“displaced”), and agents who are given higher status seats in Y than in X (“displacer”).

For any allocation Y such that $\neg X \succ Y$, construct allocation Z as follows: 1) Leave all unchanged agents where they are in Y (they will prefer Z over Y as output increases), 2) move the $\frac{N-1}{2}$ worst quality displacers back to their ideal positions (they will be made worse off than they were in Y), 3) keep the highest quality excess displacers where they are (they will prefer Z over Y as output increases), and 4) meritocratically order all the displaced in the remaining seats (there are two cases: meritocratically ordering the displaced either keeps them in the same position they were in with Y and they prefer Z as output increases or it moves them up whence they prefer Z as they get more output and higher status). Hence we have constructed Z so that it beats Y .

Now consider Z versus X . First note that because X is first best, its output must exceed Z 's. Hence, (1) All those who were unchanged and all the $\frac{N-1}{2}$ displacers remain in their ideal seat and prefer X to Z due to output. (2) By construction, there is at least 1 agent (a displacer) who was not moved down in constructing Z from Y , but this implies that there is at least one other agent who lost status and prefers X to Z . Together, (1) and (2) make a minimal majority, hence $X \succ Z$.

- Proof of iii): When $\beta = 1$ all agents strictly prefer any allocation A with higher output than any allocation B . Moreover, the first best is unanimously preferred to any other allocation. Since these preferences are strict, by continuity they continue to hold for β sufficiently close to 1.

- Proof of iv): This holds trivially for $\beta = 0$ from ii). Hence consider $\beta > 0$.

Let Y denote the allocation that produces minimal output. We want to show that Y does not cover any other allocation X . We generate allocation Z such that $Z \succ Y$ but $X \succ Z$. Going from any alternative allocation X to Y results in agents whose seats remain unchanged (“unchanged”), agents whose seats are given worse status seats in Y than in X (“displaced”), and agents who are given higher status seats in Y than in X (“displacer”).

Construct allocation Z as follows: 1) Leave those who were unchanged from X to Y where they were, 2) leave the $\frac{N-1}{2}$ top-seeded displacers in Y (who are the lowest quality agents as Y is the worst output assignment) in the same seats, 3) move the remaining lowest-seeded displacers in Y (who are the highest quality agents as Y is the worst output assignment) to their original positions in X , and 4) fill the remaining seats with the displaced, maintaining their relative seat order from X .

This construction gives us $Z \succ Y$ as the unchanged gain output, the displaced gain output and weakly gain status, the top-seeded $\frac{N-1}{2}$ displacers who remained unchanged gain output, and only the (far less than majority) lower-seeded displacers who were moved down lose status and therefore might vote against Z for $\beta > 0$.

Moreover, we also get $X \succ Z$. True, the top $\frac{N-1}{2}$ seeded displacers can be better off as their status in Z exceeds what it was in X . However, the remaining displacers are moved back in Z to their original assignment in X . So they haven’t changed status. Those who never changed from X to Y , also remain in their original assignment, so they haven’t changed status either. And those who were displaced from X to Y are assigned to the remaining seats in the same order as they were in X . The key here is that since we are keeping the top $\frac{N-1}{2}$ seeded displacers where they are (and Y is the worst output allocation), they are always worse quality than the displaced and unchanged who appear below them. Hence they are all un-meritocratic moves comparing X

to Z . Therefore, output under X is less than output under Z . Hence, although the top $\frac{N-1}{2}$ seeded displacers may prefer Z to X due to their status improvement, everyone else prefers X to Z , which constitutes a majority.

ONLINE APPENDIX B. PROPOSITION 2 PROOF FROM THAKUR (2020)

This appendix reproduces, verbatim, the proof of Proposition 2 which corresponds to Theorem 2 from *Combining social choice and matching theory to understand institutional stability*, Thakur (2020), as reference to help the reader follow our proof of Proposition 3 in our paper. Since Proposition 2 deals with the $\alpha = 1$ case, the cardinal preferences can be captured by the corresponding ordinal preference ranking that maintains the preference order of seats.

Thakur (2020) writes, “Proof: There are N seats $1, \dots, N$ to be assigned to N agents $1, \dots, N$. Suppose each agent’s preference over seats is identical, say $1 \succ 2 \succ \dots \succ N$. Fixing the order of seats²² to be $1, \dots, N$. An allocation hence defines a permutation of agents over this ordered sequence of seats. For example, an allocation denoted by $5, 2, 9, \dots$ implies that agents $5, 2, 9, \dots$ were assigned seats $1, 2, 3, \dots$ respectively.

In each step, a feasible allocation movement is defined by the preferences over seats and majority rule (only allocation movements which are approved by a majority are allowed). Because we fixed the order of seats and all agents’ preferences are identical, by definition, the composition of feasible allocation movements is also feasible. This is because it only matters which positions are moved, rather than who is in those positions. The set S of permutations defined by the feasible reassignments is thus closed under composition and thus a subgroup of S_N (the finite symmetric group on N objects) because for any element $s \in S$, it has a *finite* order k in S_N . Thus s^k and s^{k-1} , which are the identity and inverse of s respectively, are in S because of closure under composition. We want to show that group S is the full group S_N .

An N -cycle is the movement from $1, 2, \dots, N$ to $N, 1, 2, \dots, N - 1$. A $1 : 2$ transposition is the movement from $1, 2, \dots, N - 1, N$ to $2, 1, 3, \dots, N$. That the set of all possible permutations, S_N , is generated (i.e., a successive composition

²²This is key, if you fix the order of agents and try permutations of seats on this order, it does not form a group, as movements cannot be composed. Namely, the product of two movements which are individually majority-approved path from some initial allocation, cannot necessarily be combined to a majority-approved path (namely, the first movement leads to a new status quo, and hence the second movement might not be majority approved from this new status quo).

of these two operators) using these two movements, is a theorem in abstract algebra for finite symmetric groups, stated usually as $(1, 2)(1, 2, \dots, N)$ generate the group S_N .

In this setting, an N -cycle is allowed by $N-1:1$ vote. Moreover, $1 : 2$ transposition is shown for odd $N \geq 5$ and even $N \geq 6$ with a sequence of majority approved movements, in Figure 1²³. (While N -cycles work for $N \geq 3$ by majority rule, for $N = 3$ and $N = 4$, it is not possible to get any consecutive transposition like $1 : 2$ or $N : 1$ by majority rule).

Figure 1. Sequence of allocations for 1:2 transposition for N odd (left) and N even (right).

Fixing the sequence of seats $1, \dots, N$, the figures show the sequences of majority-approved allocations $A_0, A_1, \dots$ that lead to a $1 : 2$ transpositions for odd N (left) and even N (right). For odd N , at each step in the sequence from $A_1, \dots, A_N$, the bottom $\frac{N+1}{2}$ agents to the top, hence approved by $\frac{N+1}{2}$ agents. For even N , at each step in the sequence from $A_1, \dots, A_{\frac{N}{2}}$, the bottom $N - 2$ agents are moved to the top, hence approved by $N - 2$ agents.

Seat	A_0	A_1	A_2	A_N
1	1	$\frac{N+1}{2}$	N	2
2	2	$\frac{N+3}{2}$	2	1
...	...	...	1	3
...	...	...	3	4
$\frac{N-1}{2}$	$\frac{N-1}{2}$	N-1	...	...
$\frac{N+1}{2}$	$\frac{N+1}{2}$	N	$\frac{N-1}{2}$	...
$\frac{N+3}{2}$	$\frac{N+3}{2}$	2	$\frac{N+1}{2}$	...
...	...	1	$\frac{N+3}{2}$	...
...	...	3	...	...
N-1	N-1	...	...	N-1
N	N	$\frac{N-1}{2}$	N-1	N

Seat	A_0	A_1	A_2	$A_{N/2}$
1	1	3	5	2
2	2	4	6	1
3	3	...	...	3
4	4	...	...	4
...	...	...	N-1	...
...	...	...	N	...
...	...	N-1	2	...
...	...	N	1	...
N-1	N-1	2	3	N-1
N	N	1	4	N

”

²³These paths can be generated more efficiently, but for simplicity of the diagrams, we have illustrated paths with N and $\frac{N}{2}$ steps.