ECONtribute
MARKETS & PUBLIC POLICY

ECONtribute
Discussion Paper No. 084

Endogenous Organizational Restructuring: Status, Productivity, & Meritocratic Dynamics

Ashutosh Thakur Jonathan Bendor

April 2021 www.econtribute.de

UNIVERSITÄT BONN Universität zu Köln

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany´s Excellence Strategy – EXC 2126/1– 390838866 is gratefully acknowledged.

ENDOGENOUS ORGANIZATIONAL RESTRUCTURING: STATUS, PRODUCTIVITY & MERITOCRATIC DYNAMICS*

ASHUTOSH THAKUR, JONATHAN BENDOR

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 own status in the organization, the organizational output, their friends' welfare, and the quality of their own departmental colleagues. Meritocratic assignments are reached with probability one when agents value organizational output even with epsilon weight, provided friend networks and departments are not too large. We also characterize the effects of various voting rules, agendas, and specializations on the paths and the stability of organizational structures.

Date: April 27, 2021.
* A special thanks to Steven Callander, Alvin Roth, Paulo Somaini, Ken Shotts and participants of Political Economy Seminar at Stanford and 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.

Ashutosh Thakur, Stanford Graduate School of Business and University of Cologne, adthakur@stanford.edu
Jonathan Bendor, Stanford Graduate School of Business, jonathan.bendor@stanford.edu.

1. Introduction

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 departments? We examine the dynamics of organizational personnel allocation when those who are being assigned positions can themselves lobby to determine who gets what position. Whether it is politicians being assigned to committees in a legislature (Thakur, 2017), workers being assigned to positions in a firm (Cowgill and Koning, 2018), bureaucrats being assigned to jobs in the civil service (Thakur, 2019), or cadets being assigned to various branches of the army (Sonmez and Switzer, 2013), internal labor market changes can be extremely consequential for the satisfaction of the organization's members, the cohesiveness of its teams, the longevity of its departments, and the efficiency of collective production.

A desire to restructure an organization arises endogenously as members seek to advance their own interests. A member of the organization has preferences over being assigned to different positions and may be envious of other members' positions. This envy may arise from a desire for career advancement, a better work-life balance, more interesting work, increased power or responsibility, or simply a position better suited to one's qualifications and experience. Apart from the preferences over one's own seat assignment (a private good), individuals may also value the organization's output (a collective good), the well-being of friends (an agent network-based good), and have preferences concerning the composition of their team (a position network-based good). Consider these in order. First, individual members might value the organization's collective output for a variety of reasons: workers may have stock-based compensation directly reflecting their company's profitability and productivity (Hannes, 2006); politicians might want to pass legislation, partly to advance their policy agenda but also to claim credit, thereby boosting their re-election chances (Mayhew, 1974); an army's combat effectiveness can be a matter of life or death for a cadet. Second, 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 people in their in-group at the expense of those in their out-group. Third, having more capable individuals in one's own team or department can enhance group-expertise — e.g., politicians with experience in certain policy areas can, if assigned to relevant legislative committees, help those committees make better-informed choices (Krehbiel, 1992) — or increase bonuses from incentive pay schemes tied to team output (Boning et al., 2007).

Organizations are thus being pushed and pulled to reorganize in many different directions. Which forces succeed depends on which coalitions can arise endogenously in pursuit of their shared goal. In order to secure its desired restructuring, a coalition must be large and powerful enough to defeat competing coalitions and to lobby influential third parties. For example, in wartime lives are at stake and the army's operational efficacy is paramount; consequently, the common desire of army officers to enhance overall combat effectiveness might induce nearly all of them to support meritocratic re-assignments of officers to positions. On the other hand, in a family firm, a kin-based coalition might block the promotion of a talented outsider into upper level management and retain less productive dynastic family CEOs who don't maximize profits (Perez-Gonzalez, 2006). Relatedly, coalitions based on cliques or high-solidarity groups based on race, gender, or similar attributes can block certain promotions or assignments and consequently generate segregation, discrimination, and glass-ceiling effects within the workplace (Smith, 2002).

Sometimes this pulling and pushing toward different restructuring trajectories involves not ascriptive properties (kinship, ethnicity, gender) but instead the structural features of modern organizations themselves: powerful departments can impair the functioning of smaller units. Examples include Xerox's copier division which dominated Xerox's PARC division — though the latter had invented one of the earliest personal computers, it failed to monetize its innovation (Smith and Alexander, 1988) — and Imperial Germany's army which dominated Germany's diplomatic corps and, ultimately, the entire government (Kitchen, 1975), to the country's (and Europe's) detriment.

Which of these endogenous coalitions succeeds in pushing their favored reorganizations and consequently how people in firms and public agencies are assigned to positions often depends on a complex mix of informal as well as formal procedures. In this paper we simplify this complex mix by representing it as a formal process — that of voting. This simplification enhances tractability and lets us tap voting theory's analytical power to generate crisp results.[1] Further, although voting is an idealization of often informal pressure on an organization to alter its status hierarchy or otherwise restructure, sometimes formal voting is literally used: e.g., when political party members vote on alternative committee assignment mechanisms and change allocations in a legislature[2] or when employees of a company who are compensated in stocks with voting rights propose structural and personnel changes.[3]

[1]The crispness derives partly from the deterministic property of standard voting theory: either a coalition is winning or it is not. Contest functions, which represent the probability that one side beats another, can also be used to represent informal influence processes in organizations but tend to yield less crisp results. Here we opt for one side of this trade-off.

[2]For example, in anticipation of the impeachment hearings for President Trump, "some GOP members 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

[3]Consider, for example, employee activism sparred by significant share-holding with voting rights of Amazon employees via their compensation. "Amazon employees are using their company-issued stock to pressure top

The path of restructuring traversed by the organization thus depends on how members' preferences over the aforementioned goals give rise to coalitions which push for changes that advance their members' interests and how organizations resolve these conflicting interests. The possibility of endogenous organizational restructuring raises several new questions. Which allocations are stable? Are stable allocations output-efficient and are they meritocratic in their assignment of talent to positions? On what path does the organization traverse and are stable allocations reached? How do these dynamics and stable allocations respond to changes in i) the process of influence, as represented by a (possibly weighted) voting rule, ii) the protocol governing what challenges to the status quo are admissible (i.e., the voting agenda), and iii) the weights individuals assign to status, organizational output, friendship networks, and team/department productivity?

Our main finding, Theorem 1, is that even when there is maximal rivalry and competition amongst agents because of correlated preferences over seats (e.g., a commonly held status ordering), if individuals even slightly value organizational output then organizations move to the first-best allocation where agents are assigned meritocratically, i.e., based on quality so as to maximize organizational output. Importantly, for individuals whose seats assignments are not changed by the proposed alternative allocation, and who thereby might be indifferent to the proposed change, valuing organizational output even only very slightly induces them to support output-improving changes and oppose those that would reduce output. It is this otherwise indifferent mass of abstainers who join the pro-change coalition and overwhelmingly push towards meritocracy-enhancing and output-increasing reorganization.

This disciplining effect of valuing organizational output carries through even in the presence of friends who look out for each other's interests, as long as friendship networks are not very dense and no agent is extremely well-connected. The failure of these provisos to hold can explain why family firms can be inefficient and non-meritocratic and why glass-ceiling effects occur when certain groups (male or big racial/ethnic groups) are disproportionately large within an organization, as well as highlighting how bigger organizations might have fewer inefficiencies because social ties and networks are less able to block output-increasing reassignments. The disciplining effect of valuing organizational output also carries through in the presence of individuals wanting higher quality members in their own teams/departments, as long as these teams are not so large as to unilaterally monopolize the restructuring process.

We further show that if there are advantages to specialization (Ferguson and Hasan, 2013), via match-based output efficiencies, then an organization may fail to reach the global first-best because it gets stuck at a local maximum. Moreover, we find that the organizational output effect can overshadow individual preference even if preferences are not fully correlated.

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/

For the most part we focus on voting by majority rule for ease of exposition. Our main result from Theorem 1, is robust to many voting rules where power is decentralized, such as q-rules (e.g., super-majority rules to change political party rules or peer-reviews within firms). More precisely, Appendix A establishes that, when representing voting rules in terms of decisive coalitions, Theorem 1 is robust to any voting rule where no agent is so powerful that he belongs to every decisive coalition. (In social choice terms, agents who belong to every decisive coalition are collectively called a collegium. So Theorem 1 is robust provided no collegium exists.) The voting rule can even allow for an agent's voting power to be endogenously affected by the position that the agent happens to occupy at the time of a particular vote. This robustness illustrated in Appendix A makes the interpretation of our voting procedures less stylized and closer to informal collective decision-making through debate, negotiation, and internal power play. We also analyze the effect on the path of organizational restructuring when voting rules admit a collegium arising from sharply different influence structures such as unanimous consent, an oligarchy (e.g., board of directors), a dictatorship (e.g., CEO), and an agent with veto power (e.g., a large shareholder). The more nuanced dynamics in these cases result in meritocratic assignments for all non-collegium members and potentially non-meritocratic assignments for members of the collegium. Such influence structures can thus lower organizational output.

Related Literature. While Thakur (2020) considers the static problem of existence of institutionally stable allocations, our paper analyzes the dynamic paths of allocations taken by the organization and introduces preferences for organizational output, friendship networks, and the quality of departments in addition to rank order preferences over one's own assigned seat. There is also a growing literature in computer science on Popular matching, which focuses on stability using plurality rule (see review article by Manlove (2013) and Cseh (2017)). This literature deals with the plurality rule as the choice rule in discriminating between allocations and considers the existence, characterization, and complexity of finding a Condorcet winner given one-sided (Abraham et al., 2007; Sng and Manlove, 2010; Manlove, 2013; McDermid and Irving, 2011; Kavitha and Nasre, 2009; McCutchen, 2008; Kavitha et al., 2011) and two-sided voting in bipartite graphs (Huang and Kavitha, 2013; Kavitha, 2014; Cseh et al., 2017; Cseh and Kavita, 2016) and non-bipartite graphs (Chung, 2000; Biro et al., 2010; Huang and Kavitha, 2017). In addition to analyzing plurality rule we examine many other voting rules and in addition to rank order preferences we study several other empirically common motives that affect individuals' preferences. Lastly, analyzing the stability of the institutional choice itself shares the spirit of self-stable voting rules and constitutional choice literature (Barbera and Jackson, 2004; Messner and Polborn, 2004; Maggi and Morelli, 2006; Koray, 2000).

2. Model

Consider the problem of assigning $N \geq 3$ agents (denoted by $i \in \{1, ..., N\}$) to N seats (denoted by $s \in \{1, ..., N\}$) within an organization. For ease of notation, assume N is odd, unless otherwise stated. This is a balanced market with an allocation being a one-to-one matching between agents and seats.

Each agent i has a latent 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 $Y(A)$. Denote by $A(i)$ the seat that i is assigned in allocation A. We assume that the output production function $Y(\cdot)$ 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: $Y(A') > Y(A)$. This means that the marginal gain from having a higher quality agent is more in a higher status seat.[4]

Assume that there is a perfectly agreed-upon status hierarchy over seats, i.e., all agents have identical preferences over seats[5] in that $u_i(s) > u_i(s')$ if $s < s'$ for all i.[6] This is relaxed in Section 7.

The utility of agent i from allocation A is

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

where $s(A)$ is the agent who is assigned seat s in allocation A, weights $\alpha + \beta + \gamma + \zeta = 1$, $Friend_i$ denotes the set of agents who are friends with i, $Team_i$ denotes the set of seats which also belong to i's team, and f_i are (because higher skill is represented by lower numbers) decreasing functions. The first term is the agent's payoff from the status of his position. The second term is the payoff from the performance of the organization as a whole.[7] The third term is the payoff that agent i derives from the status enjoyed by i's friends (i.e., from

[4]Increasing 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 $s < s'$, $Y(A(i) = s) - Y(A(j) = s) > Y(A(i) = s') - Y(A(j) = s')$ holding all other agents' assignments in A constant.

[5]From 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. See Shubik (1971) for an early analysis of games of status.

[6]We 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. It would be interesting to explore endogenous salaries, e.g., as a function of organizational output or as side-payments to compensate people for doing organizationally valuable work of low prestige; we leave these topics for future work.

[7]Models in which the outputs of different departments are measured on different dimensions (e.g., marketing department cares about sales, finance department cares about profit) are left as future work. The present model considers a single dimension of organizational output.

their assigned seats).[8] The fourth term is the payoff from the quality of those assigned to your team: having higher quality agents on your team can enhance the group's expertise or increase bonuses from team output-based incentive schemes.[9]

Agent i votes over two alternative allocations if and only if he has a strict preference over these allocations; otherwise he abstains. The voting rule for the organization's allocation choice is majority rule, namely, alternative allocation A' replaces A if and only if a majority of agents strictly prefer A' over A. We consider other common voting rules in Section 8 and Appendix A. We assume voting is myopic.

Voting can of course be an idealization of the endogenous and often informal pressure on the organization to undergo restructuring and reorganization. Certain organizations literally use formal voting to change institutional procedures, e.g., when party members vote on alternative committee assignment mechanisms and change allocations in the legislature or when employees of a company who are compensated in stocks with voting rights propose structural and personnel changes. Moreover, subjective peer-evaluations in the form of performance appraisal reports or 360 degree evaluations from managers, peers, and subordinates often plays an integral role in determining promotions and organizational restructuring. In other instances, group decisions are reached organically through debate, discussion, and deliberation. In such applications, we interpret voting to be an idealization of providing an opinion and arguing in favor of or against a proposed change. 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 the important feature that opinions coming from certain agents or from agents in certain positions can carry different weight and influence. 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 the 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 from only that employee's managers and workers working directly under that employee's supervision. Organizations discipline their intra-organizational decision-making in different ways; we are trying to capture this richness in a stylized manner through voting.

[8]The same mathematical formulation can account for kin, professional contacts, or any other individual-based social tie. Note that for now we are only considering first-order friends and not network effects such as friends of friends.

[9]With 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.

3. The Consequences of Valuing Organizational Output.

We start the analysis with agents considering only the utility they get from their own seat assignment; we then consider the incremental effect of agents also valuing organizational output.

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

> **Proof**: For any candidate allocation, a coalition of size $N-1$ can always be made better off by moving the agent who was assigned the highest status seat down to the bottom and moving all $N-1$ others up one position in the status hierarchy.

The shared status ranking of positions in the organization implies that preferences are fully correlated. Hence, no allocation is majority stable 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 allocations, given perfectly correlated preferences over seats and the ensuing rivalrous nature of the private-status-only environment.

We now turn to examining how the organization evolves dynamically and incrementally. To do this, we restrict the alternative allocation protocol to *swaps*: allocation A_0 is voted against an alternative allocation A' in which two agents i and j's seats are swapped and all other assignments remain fixed. We allow for any alternative allocation protocol where for any given allocation as the status quo, all potential allocations which are 1 swap away are not always excluded, dynamically or in a probabilistic sense.[10] We discuss other alternative allocation protocols in Section 8.

Lemma 2. *If $\alpha = 1$ and the alternative allocation protocol is restricted to swaps then any allocation is stable.*

> **Proof**: Any swap leaves all $N-2$ *uninvolved* agents (those agents whose allocations are not changed by the swap) indifferent and makes the agent who loses status worse off while making the agent who gains status better off. Hence, any alternative allocation loses by a vote of 1 to 1, with $N-2$ abstaining.

[10] This captures a rather broad class of alternative allocation protocols including i) random swaps proposed as alternative allocations, ii) random (non-zero) recognition probabilities for each agent to be a swap-proposer, iii) rotating rosters of swap-proposing power across agents, and iv) rotating rosters of swap-proposing power with consecutive swap-proposing power for the same agent if swaps are successfully passed. Alternative allocation protocols which are excluded from this class are those where the proposing power remains concentrated locally amongst a subset of agents, e.g., dictatorial proposal power or rotating roster or random recognition power restricted to a small subset of the agents.

If agents only care about their own allocation then restricting to swaps generates overwhelming stability because the involved pair of agents deadlock, all agents whose seats are unchanged remain indifferent and hence abstain from voting, so the change falls far short of majority approval.

This, however, is a knife-edge result. Even epsilon weight on organizational output changes things dramatically.

Theorem 1. *If* $\alpha+\beta = 1$, $\alpha > 0$, $\beta > 0$ *and the alternative allocation protocol is restricted to swaps then i) the unique stable allocation is the first-best and ii) the first-best is dynamically reached with probability 1.*

> **Proof**: i) The meritocratic first-best allocation is stable since any swap reduces output, thereby making all $N-2$ uninvolved agents worse off. There is no other stable allocation because every other allocation has at least one inverted pair such that $i < j$ despite $A(i) > A(j)$ and swapping such pairs makes $N-2$ uninvolved agents better off. ii) Corollary 1.1 along with the finite set of possible matchings implies the first-best is reached with probability 1.

When agents also value organizational output, *uninvolved agents*, whose own assignment does not change, vote for (against) output-improving (-decreasing) swaps. All output-improving swaps pass with an overwhelming $N-1$ to 1 vote, which has three implications. First, for majority rule, $N-1$ voting in favor of output-improving swaps means that there is slack. To illustrate this, notice that instead of swaps, any output-improving change in allocation that changes at most $\frac{N+1}{2}$ agents' seats, will be passed under majority rule. This is because at least one agent is assigned a higher status seat in the new allocation, and this agent(s) along with the $\frac{N-1}{2}$ uninvolved agents constitute a majority. Second, the result is robust to plurality rule, to any q-rule that is not unanimity rule, and as long as no agent(s) has veto power (see Section 8). Lastly, as is often the case in practice, different seats or different agents' opinions can carry different weights. Even with such weighted voting, as long as no single agent's weight is enough to be decisive by himself, the convergence to the first-best is guaranteed. Appendix A provides necessary and sufficient conditions for convergence to the first-best under a wide array of voting rules.

It is not surprising that the first-best outcome is obtained when β is large; it is $\beta = \epsilon$ that makes this a striking result. As the rest of the paper will demonstrate, the overwhelming impact of the otherwise indifferent $N-2$ agents, their involvement triggered by $\beta > 0$, is a robust effect. Giving employees[11] and executives[12] stock-based compensation has been hotly debated yet frequently used in practice with the intention of aligning incentives and inducing effort. Theorem 1 illustrates the effects of such policies on the endogenous pressure for reorganizing personnel and talent within the organization.

Corollary 1.1. *If $\alpha+\beta = 1$, $\alpha > 0$, $\beta > 0$ and the alternative allocation protocol is restricted to swaps then output strictly increases over the dynamic path.*[13]

> **Proof**: If output decreased then all $N - 2$ uninvolved parties would vote against the proposed swap; if output remained constant then all $N-2$ uninvolved parties would be indifferent and abstain. Hence the only swaps that pass are output-improving changes where the uninvolved parties support the new allocation.

This dynamic property has an interesting implication in settings in which 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 date t can reduce organizational output. From that time forward, however, the corollary implies that output will once again start increasing. The logic of the $N - 2$ uninvolved agents simply restarts, hill-climbing on output gradients that have been reset by the recent shock.

Corollary 1.2. *If $\alpha + \beta = 1$, $\alpha > 0$, $\beta > 0$ and the alternative allocation protocol is restricted to swaps then the distance from the meritocratic ideal, $\sum_i |s_i - s_{meritocraticideal}|$, weakly decreases over the dynamic path.*

> **Proof**: Consider a pair of agents of quality $q_1 < q_2$ (lower is better) and seats ranked by status $A < B$ (lower is better). For any swap, only the involved parties have their assignments changed; hence, showing that an inverted pair (q_1 assigned B and q_2 assigned A) is weakly less meritocratic than the corresponding meritocratically ordered pair (q_1 assigned A and q_2 assigned B) can involve just the absolute deviations from the ideal for the two involved agents. It is easy to verify that $|q_2 - A| + |q_1 - B| \geq |q_1 - A| + |q_2 - B|$.

[11]Weitzman (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 lead to improved monitoring of other employees.

[12]A large theoretical literature starting from Jensen and Meckling (1976) suggests that option contracts can align executives' incentives with shareholders and induce effort. 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. Others have argued that stock options are an inefficient way to compensate executives (Jenter, 2001; Meulbroek, 2001; Hall and Murphy, 2002; Lambert and Larcker, 2002).

[13]The "dynamic path" refers to the consecutive sequence of *new* allocations reached; it omits multiple instances of the same allocation, i.e., if an alternative allocation was voted down and the allocation from period $t - 1$ is maintained in period t.

Corollaries 1.2 and 1.1 underscore the organization's drift toward meritocratic assignments and efficient production with $\beta > 0$. Instead of restricting alternative allocations to swaps, even when allowing for large changes in alternative allocation proposals, in order to move in an anti-meritocratic, anti-output direction, the proposed change must reassign a large number of agents to new seats, the status-winners must outnumber the status-losers, and sufficiently many agents must put low weight on organizational output.

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.

We have not placed any restriction on which swaps are proposed, i.e., on the exact agenda. Thus, although there could be multiple possible dynamic paths of allocations the organization could take, the $N-2$ uninvolved agents vote to ensure that only output-improving and weakly meritocracy-improving swaps are carried out to approach the stable first-best allocation.

The restriction of the agenda setter's power to manipulate the allocation (only to output-improving paths leading to the first-best allocation are majority-approved) comes from *both* the alternative allocation protocol being restricted to swaps *and* $\beta > 0$. Lemma 2 illustrates that restricting the alternative allocation protocol to swaps alone leads to all allocations being stable. On the other hand, the next two propositions illustrate the unconstrained power of the agenda setter when the alternative allocation protocol is unrestricted. Even with the incremental addition of small $\beta > 0$ (Proposition 2), the agenda setter retains unconstrained power as in the $\alpha = 1$ case (Proposition 1, from Thakur (2020)) to reach any alternative allocation from any initial allocation, akin to McKelvey's Chaos Theorem (McKelvey, 1976). That is, there exists a majority-approved sequence of reassignments from any allocation to any other allocation for $N \geq 5$.

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

Proposition 2. *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 $\beta < \beta^*$ there is a majority-approved path (sequence of reassignments) from any allocation to any other allocation.*

Proof: Thakur (2020)'s proof (reproduced in Appendix C for 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, 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.

Unsurprisingly, for large enough $\beta > 0$, the agenda setter's power is restricted to organizational output-improving alternative allocations. However, Proposition 2 highlights that the key knife-edge effect of small $\beta > 0$ in Theorem 1 is caused jointly by the restriction of the alternative allocation protocol to swaps and $\beta > 0$. Hence this paper illustrates the incremental endogenous organizational restructuring as a result of agents valuing organizational output and productivity of the organization.

4. Consequences of Teams & Departments.

A nearly universal feature of organizations of even modest size is a team or departmental structure. It is thus 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 and independence to organize its own personnel, this just defines a smaller version of our problem restricted only to this sub-group of the organization. For example, personnel management autonomy is often given to franchises, subsidiaries or decentralized companies of a large transnational corporation. In this section we model teams or departments that are not given such complete autonomy in personnel allocation. The ζ term adds a pay-off to agent i based on the quality of the agents occupying the seats that constitute agent 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 brought about by informational spillovers, and improving work satisfaction due to increased competency of peers and colleagues. Sometimes, teams are given group incentive pay, making the connection to caring about the team members' qualities even more direct (Boning et al., 2007).

Proposition 3. *If $\alpha + \zeta = 1$, $\alpha > 0$, $\zeta > 0$ 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,[14] *then that largest team gets the best quality workers and they are ordered by quality within-team, while any assignment of the worse quality workers across any minority team(s) is a stable allocation.*

ii) If no team has at least a minimal majority-1 of seats,[14] *then any allocation is stable.*

Proof: We will use the shorthand "majority+1 team" to stand for "team with at least minimal majority+1 seats." i) Within the majority+1 team, swaps that increase meritocracy are passed by a majority: the only agent who votes against is the agent within the majority+1 team who is made worse off by getting a lower status seat;

[14]If 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 Y for Proposition 4).

agents in all the other teams are indifferent and abstain. Consider any swap across or within minority teams. Such swaps cannot pass because the majority+1 team and all uninvolved teams (those whose members are not being swapped) are indifferent and abstain, and the involved team getting the lower quality agent votes against the swap. Now consider what happens from a swap involving the majority+1 team and a minority team. Everyone whose assignment remains unchanged within the majority+1 department votes for a switch which improves the quality in their unit, the minority team votes against being assigned a worse agent, and all uninvolved teams abstain since they are indifferent. This is always a majority since the largest department has at least one more agent than the minimal majority. ii) Any swap within teams only has the involved team voting and that never constitutes a majority. And any swap across teams has the benefiting team voting for, the losing team voting against, and all uninvolved teams abstaining, which does not form a majority even if both agents being swapped vote in favor of the swap.

The logic of uninvolved agents extends to *uninvolved teams*: teams whose agents are not changed under the alternative allocation remain indifferent and abstain. Hence, unless there is a majority team, all allocations are stable. If, however, a majority team exists then, because it votes to add the highest quality agents to itself, it may skew meritocratic assortment and produce inefficient organizational output.

Proposition 4. *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,[14] *then stable allocations have every team ordered by quality within-team. However, depending on the parameter weights there may be agents in the majority+1 team who are under-placed.*

ii) If there is no team with at least minimal majority-1 seats,[14] *then with probability 1 the first-best is reached via output-improving swaps.*

Proof: We will use the shorthand "majority+1 team" to stand for "team with at least minimal majority+1 seats." i) First, consider a within-team swap. Only those that are output-improving will pass; the $N-2$ uninvolved agents vote for such changes. Second, consider a swap across two minority teams. If the organizational output is decreased, the uninvolved majority+1 team will vote against; if output is increased, the uninvolved majority+1 team will vote for. Hence only swaps that increase output are passed. Lastly, consider an across-team 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 Y, and functions u_i and f_i. There are a few cases to consider: a)

The majority+1 team will always vote for a swap that makes quality of its own team better off and increases output. Hence such swaps always pass. b) The majority+1 team will never vote for a swap that makes the quality of its team worse and decreases output. Hence such swaps never pass. c) A swap that decreases majority+1 team quality and increases output only passes 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 the majority+1 team 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. ii) First, consider a within-team swap. Only those that improve output will pass, with the $N-2$ uninvolved agents voting for. Second, consider a swap across 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 worse 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.

When agents care about organizational output, uninvolved teams again vote for all output-improving swaps and vote against any 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[15] can potentially skew meritocratic assortment and lead to sub-optimal organizational output by voting to attract higher quality agents who would have served the organization better in another team.

These results highlight the potential downside effects of having very asymmetrically sized groups, teams, or departments within organizations. By monopolizing the vote over organizational changes, very large teams can mis-allocate talent and produce sub-optimal output. For example, the dominance of the Xerox copier division is accredited 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/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 pivotal changes that necessitate internal re-organizations.

[15]Myopic voting implies that on some paths, some majority team agent i might vote to have a better quality team-mate join his unit but as a result i is later demoted within the majority team or ultimately kicked off it altogether.

5. 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 the 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 boy's networks based on discrimination and segregation on gender, race, ethnicity, etc (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, so an individual is indifferent with regards to his friends term in the utility, when both agents who are being swapped are his friends. 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. We later comment on the implications for organizational structure of assuming that friendship is symmetric and that groups of friends are fully connected. Of course, enmity or hostility is functionally equivalent to being friends with the other agent.

Proposition 5. *Assume that $\alpha + \gamma = 1$, $\alpha, \gamma > 0$ and the alternative allocation protocol is restricted to swaps. If γ is small enough*[16] *such that any improvement in status of i causes i to vote for swap improving himself then the only swaps that are carried out are those where i improves his status at expense of j if and only if i's friends who are not j and not also j's friends constitute majority-1.*

Proof: Any swap improving the status of i and worsening that of j leaves those who are friends with both indifferent, those who are friends with neither indifferent, those who are only friends with i better off, those who are only friends with j worse off, j worse off, and i better off. Hence there are three possibilities:

(1) If i plus those who are only friends with i constitute a majority then any such swap is carried out by majority.

[16] If γ is large enough such that i might vote against his own status for the benefit of his friend j's status, what changes is that we now get someone who is altruistic or selfless, and hence the two agents i and j who are involved in a swap might themselves vote against/for the swap.

(2) If j plus those who are only friends with j constitute a majority then no such swap is carried out by majority.
(3) If neither of the groups above are in majority then no such swap is carried out.

Depending on the friendship network, stable allocations can be far from meritocratic and hence organizational output can be sub-optimal. Only those who have many connections can gain status through the support of their network.

Proposition 6. *If $\alpha + \gamma + \beta = 1$, $\alpha, \gamma, \beta > 0$ and the alternative allocation protocol is restricted to swaps then the set of stable allocations has weakly more output and is weakly more meritocratic relative to $\beta = 0$.*

Proof: Given that $\beta > 0$, some of the allocations that are stable in the $\beta = 0$ case are unstable here because those who are friends with both and those who are friends with neither of the swapped agents i and j all vote in favor of output/meritocracy-improving swaps. Moreover, depending on the weights of the parameters, those who are friends with just one of i or j might also now only favor meritocracy-improving swaps. So the set of stable allocations is weakly more meritocratic and has weakly higher output than the $\beta = 0$ set.

To illustrate the disciplining effect of agents valuing organizational output, consider a simple example with $N = 5$. Suppose in the initial allocation agent i is assigned to seat i (denoted by the agents' permutation 12345), agents 1, 4, and 5 are friendless, and agents 4 and 5 are each friends of both agents 2 and 3. Under $\beta = 0$, there are multiple stable allocations depending on the agenda. If agent 2 or 3 vies first for seat 1, the stable allocation reached is either 23145 or 32145 respectively. Moreover, for $\beta > 0$ small enough so that friendship still overpowers any output considerations, both 23145 and 32145 can still be stable. However, 32145 could be unstable for intermediate small β as output considerations allow agent 2 to beat agent 3 for seat 1. And for large $\beta > 0$, it could be that only 12345 is stable. (This also shows that the set of stable allocations under $\beta > 0$ may be an improper subset of the set of stable allocations under $\beta = 0$ or may also have other allocations not found in the $\beta = 0$ set.)

Propositions 5 and 6 indicate that friendship and other informal connections can help an individual climb an organization's status hierarchy. Luthans, et al. (1988) found the most successful managers to spend 70 percent more time engaged in networking and 10 percent more time in routine communication activities compared to their counterparts. Our results highlight that a 'useful' friend is one who is not friends with many others: such agents are less likely to be friends with both of the parties involved in a swap and thus avoid facing a conflict of friendship that, under our assumption, causes them to abstain due to indifference. If very well networked i's friends are a majority-1 of the organization, i might get to the highest-status position. If there are multiple agents who have at least a majority-1 friends, none can displace another. In this case the agenda order of swaps matters (i.e., if i can propose a swap first, i might get the highest-status position) as illustrated in the example above.

If, however, agents also cared about organizational output then distortions arising from friendship ties can be alleviated. This explains 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).

Now consider groups of agents who are friends with everyone in the group: e.g., tight friendship cliques or high-solidarity groups based on race, gender, ethnicity, or the like. Such networks can support tiered organizations: e.g., a male-dominated organization in which men are the majority and hence all enjoy 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.

Finally, if the friend networks in an organization are dense and overlapping then person-to-position assignments can remain far from meritocratic and output can be far from optimal. This is consistent with the ample literature documenting inefficiencies in family firms.[17]

6. Specializations & Expertise

Certain positions are associated with particular specializations, where skill set match can arise from education qualifications or work experience. Assigning an agent with the right skill set to the right position can increase productivity and specializations can affect career trajectories within an organization (Ferguson and Hasan, 2013). In this section we explore the effects of introducing specializations in to our model.

Assume that there are $m = 1, ..., M$ types of specializations or skill sets. Each seat requires a certain type of specialization; each agent has a specialty. Mismatches reduce output: if agent with speciality p is assigned to seat of speciality q then there is a loss of output $d(p, q)$. Thus, the mismatch penalty can differ based on the distance between the specialities of agent and seat, as measured in some specialization space.

Lemma 3. *Assume $\alpha + \beta = 1$ and $\beta > 0$. For any arbitrary $d(p, q)$ penalties there does not necessarily exist a path of swaps passed by majority rule from an arbitrary allocation A_0 to the first-best allocation A^*.*

[17]Many 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).

Proof: For a counterexample to the claim that such a path always exists, suppose there are 3 players of types $\{a, b, c\}$ and 3 seats of types $\{A, B, C\}$, where like alphabet types denote a specialization match. Suppose the first-best $A*$ matches specialities: aA, bB, cC. Moreover, make penalties and outputs such that the uniquely second-best output allocation is cA, bC, aB. There is no direct swap from the second-best allocation to the first-best and by design, any swap away from the second-best decreases output and is hence voted down by the uninvolved agent and the agent whose status would be decreased by the proposed swap.

Interestingly, the analogue of this lemma still holds true even if we simplify the structure of specialization penalties to be such that any mismatched agent-seat pair generates a fixed output loss $d > 0$ (while $d = 0$ for matched agent-seat pairs). The same proof-logic — an example of the second-best not having a direct swap leading to the first-best — works here as well. We present a numerical example for exposition and clarity.

Example 1. *Consider 3 agents i, j, k ordered by quality $q_k > q_i > q_j$ who are of types I, J, I. Assume that the seats are ordered by status by J, I, I. We posit the meritocratic output Y for each possible matching and the adjusted output Y' which includes specialization mismatch penalties $d = .49$.*

Matching	Y	*penalty*	Y'
i, j, k	*4*	$-2d$	*3.02*
j, i, k	*3*	$-0d$	*3*
j, k, i	*4.2*	$-0d$	*4.2*
i, k, j	*5*	$-2d$	*4.02*
k, i, j	*5.1*	$-2d$	*4.12*
k, j, i	*4.3*	$-2d$	*3.32*

The allocation j, k, i is the first-best with $Y' = 4.2$. The second-best is k, i, j with $Y' = 4.12$; however, there are no swaps from this allocation directly to the first-best. Hence there is no path of output-increasing swaps from the second-best to the first-best allocation that would have passed with majority rule.

Adding specialization of types across seats and agents introduces added dimensions, thereby triggering a curse of dimensionality. This kicks in especially when proposed assignment changes are restricted to pairwise swaps. This property together with myopic voting can get an organization hung up on 'local' maxima.

7. No Common Status/Preference Ordering

We now consider relaxing the assumption that there is a shared preference ordering over seats. People could have different preferences over positions: 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. When the alternative allocation protocol is unrestricted, 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, etc. to unravel majority stability. We consider what happens when agents value organizational output.

Lemma 4. *For any arbitrary preferences of agents over seats, if $\alpha+\beta=1$, $\alpha>0$, $\beta>0$ and the alternative allocation protocol is restricted to swaps then the output-maximizing allocation is the unique stable allocation and it is reached with probability 1 along an output-improving path.*

> **Proof**: For any non output-optimal allocation, $N-2$ uninvolved agents vote in favor of (against) any output-improving (-decreasing) swap. The output-optimal allocation is reached with probability 1 since there are a finite set of possible matchings. The output-optimal allocation is stable because any swap decreases output and hence is voted against by the $N-2$ uninvolved agents.

The logic of $N-2$ uninvolved agents voting for output-improving swaps carries through even here, thereby leading to the allocation that maximizes output. However, many agents could be misplaced in terms of their own preferences over seats: the $N-2$ uninvolved agents might force an involved agent into a position they despise. Hence the allocation that maximizes output might differ from the social planner's first-best, which maximizes the (weighted)-sum of everyone's utilities. The social planner's first-best balances output and positional preferences, but output maximization is so strong in the voting rules 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): the obsession with output does arrange agents by quality and thus maximizes output, but this comes at the expense of agents' own preferences over positions.

8. Other Common Voting Rules

Here we consider what happens under other common voting rules.

First, the generalization of majority rule to a q-rule, $\frac{N+1}{2}<q<N$ changes none of the fundamental results. The underlying logic of the $N-2$ uninvolved agents voting only for output-improving swaps remains in force.

Second, under 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 large effects of $\beta=\epsilon$ given q-rules disappear. Moreover, even with friends (unless i is friends with all other $N-1$ agents) and departments, a very high β weight is required to get any movement. Thus, unanimity rule highlights the fundamental rivalry inherent in choices over positions and status.

Third, consider using the dictatorial rule where 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 output loss might be more than his status gain. That is, the dictator puts himself at the highest seat such that if he were to increase to higher status position, $-\beta * \Delta Y > \alpha * \Delta status$. Thus, 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. With regards to 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 which improves output but makes the allocation farther in terms of swap distance from the meritocratic ideal.[18]

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 anything that decreases either his status or organizational output. So once again, stable allocations are meritocratic other than in i's own position. Moreover, if i starts off under-placed, he will be moved up to his ideal position because all $N-2$ uninvolved agents vote for output-improving swaps. But if he is ever over-placed, either initially or as a result of a swap, 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 i's support protects an agent against being moved down the status hierarchy. And being in the department of a veto player ensures that one's unit can never be made worse off even if it is a minority department.

[18]Suppose 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 and how such instruments can create incentives to act against long-run growth via short-run manipulations.

Fifth, consider an oligarchy defined as a set of agents G such that all $i \in G$ have a veto and are collectively decisive: if all $i \in G$ prefer $A' \succ_i A$ then A' is passed. If $\alpha = 1$ then everything is stable. 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 starting 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 — hence he makes himself as well off as possible (his optimal position depends on the β weight) and orders everyone else meritocratically — whereas oligarchs are rivalrous against each other. They all vote to impose meritocratic discipline on everyone else, but they do not want to suffer such discipline themselves and they can 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. 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.

Hence, if a shock to the organization's task environment changes the agents' quality ordering (per the discussion following the Theorem), an oligarchy that had previously been consistent with meritocracy can block adjustments that would move the organization to the new post-shock meritocratic ordering. Newly over-placed oligarchs use their authority to block such changes.

Next, consider the organization using a plurality rule. The uninvolved $N-2$ voting logic 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 moves, 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 $N-2$ uninvolved 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 $\beta = 0$.

Proposition 7. *Under plurality rule, if $\alpha + \zeta = 1$, $\alpha, \gamma > 0$ and the alternative allocation protocol is restricted to swaps 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 any path converges with probability 1 to one of these stable allocations.*

> **Proof**: Plurality rule implies that agents who have more friends outvote those with fewer friends since uninvolved agents who are friends with neither of the swapped agents and agents who are friends with both swapped agents abstain. Hence we reach an ordering of agents into seats in order of status based on the absolute number of friends they have. For a set of agents who all have the same number of friends, any assignment of their seats is stable. Depending on the path of swaps proposed, one of these stable allocations is reached with probability 1 since there are a finite number of possible matches and the only swaps that are passed improve the status of one agent with more friends at the expense of another agent with fewer friends.

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, within a friendship bloc every allocation is stable, and any path converges to one of these strict dominance hierarchies with probability 1.

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] Namely, 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 have a say. Moreover, swaps involving different seats i, j may have different status thresholds $\underline{s}(i, j)$, i.e., a decision to replace a CEO might be restricted to the Board of Directors, whereas replacing a mid-level manager may only involve that manager's team working under him/her having some say in the decision.

Lemma 5. *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 $\alpha + \beta = 1$, $\alpha, \beta > 0$ and the alternative allocation protocol is restricted to swaps then the first-best is reached with probability 1.*

[19]Our main result of convergence to the first-best under $\alpha + \beta = 1$, $\alpha, \beta > 0$ and the alternative allocation protocol being restricted to swaps, 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 many agents having no influence or say in certain personnel changes. This underscores the overwhelming effect of the $N - 2$ uninvolved agents supporting only output-improving swaps.

Proof: If $\underline{s} = 1$, an under-qualified agent who is over-placed in seat 1 will never vote to make himself worse off for small enough β. Moreover, if $\underline{s} = 2$, then any twisted pair that involves either seats 1 and/or 2, will not pass with majority vote. For $\underline{s} \geq 3$, 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.

9. Other Alternative Allocation Protocols

Throughout most of this paper we have focused on the alternative allocation protocol to be restricted to proposing swaps. If this holds then for the case of $\alpha + \beta = 1$ and $\alpha, \beta > 0$, all convergence results follow as long as the protocol considers every feasible swap with positive probability. For example, if the alternative allocation protocol proposes a random swap then the ensuing model is Markovian with a unique absorbing state — the first-best allocation — and convergence to this output-maximizing state is guaranteed in the limit if every swap is proposed with positive probability. Moreover, the voting of the $N - 2$ uninvolved agents also leads to the output-improving and weakly meritocracy-improving path, on which consecutive allocations become increasingly stable.[20] Alternatively, one could assume that transitions from one allocation to another are governed by a contest function, e.g., transition probabilities equal to the fraction of voters preferring the alternative allocation to the status quo. Dynamics 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. Since this approach does not presume that the organization literally votes (via some formal procedure) on agent-seat assignments, it is a relaxation of the voting idealization that we have used in the present paper.

What happens when the alternative allocation protocol is not restricted to swaps? Although Propositions 1 and 2 establish McKelvey (1976)-like chaos of majority voting, adding some Markovian structure highlights the robustness of our results about organizational efficiency in the limit, given $\alpha + \beta = 1$ and $\alpha, \beta > 0$. We use the following assumptions.

A1: for any status quo allocation, the recognition probability that any agent i gets to propose the alternative proposal is positive.

[20] The notion of varying degrees of stability used in this Markovian model is related to research on differing sizes of win sets in legislative systems (Ferejohn et al., 1980; Ferejohn et al., 1984). The win set in the present context is the set of alternative allocations which would be preferred by a majority over the status quo.

A2: 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.

Theorem 2. *If assumptions A1 and A2 hold 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 space 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* $\overline{\beta} < 1$ *such that for all* $\beta > \overline{\beta}$*, the limiting distribution is degenerate, with all mass concentrated on the first best allocation; iv) the only allocation that occurs with positive probability in the limiting distribution for all* $\beta \in [0,1]$ *is the first best.*

Proof: We have a finite Markov chain with stationary transition probabilities and assumption *A2* guarantees aperiodicity. In Theorem 5i) we proved 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-preferred path of allocations from the first best) or a recurrent state (to which there exists a majority-preferred path of allocations from the first best) that belongs to the single closed recurrent class, which proves i). For β small enough, Proposition 2 implies the Markov chain is irreducible as all states communicate (there exists a majority-preferred path of allocations 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 it becomes an absorbing state as it is the Condorcet winner for β sufficiently close to 1.

Theorem 2 is not restricted to swaps; it allows for a wide array of alternative allocations which can support strategic agenda setting.[21] 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 (as it is if the alternative allocation protocol is restricted to swaps), it is the most stable allocation with the highest weight in the limiting distribution.

[21]Both assumptions *A1* and *A2* can be relaxed still further without sacrificing the analytically useful framework of a finite Markov chain with stationary transition probabilities.

10. Conclusions

In this paper we explore how the interests of agents in increasing organizational output impact the endogenous assignments of agents to positions. We find that when agents care even slightly about organizational output, due to owning shares of a company or to some intrinsic skin in the game (army placements in war or political committee assignments during legislative battles between political parties), the organization moves towards meritocratic placements and efficient output. The commonality of preference for increasing organizational output 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 unless there are extremely large teams or extremely well-connected individuals who can enlist large voting blocs to oppose meritocratic, output-improving organizational changes. 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 and reduces organizational output. Even so, 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) tend to overlook this strategic complexity.

Our model also uncovers dynamic complexities: if specialization matters then organizations can get stuck at local maxima, thereby never reaching the first-best. This is due to the curse of dimensionality: adding specializations that are imperfectly correlated with the quality dimension produces a payoff topography with local maxima for which the swaps-only protocol is ill-suited.

Importantly, caring about output adds a near-perfect $(N-1)$ commonality in preferences for organizational restructuring across individuals that stimulates output-improving swaps that lead to the first-best. 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 on the other hand the alternative allocation protocol is restricted to swaps then, because the $N-2$ uninvolved parties are indifferent, everything is stable. But even a small weight on organizational output gives parties who are not directly involved in a swap a stake in the game, thus creating a strong organizational inclination toward improving organizational output.

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] 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.

[4] Bandiera, O., Guiso, L., Prat, A. and Sadun, R. (2012), What do ceos do?, CEP Discussion Papers dp1145, Centre for Economic Performance, LSE.

[5] 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.

[6] 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.

[7] 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.

[8] Biro, Peter, Robert W. Irving, and David F. Manlove. “Popular matchings in the marriage and roommates problems.” International Conference on Algorithms and Complexity. Springer, Berlin, Heidelberg, 2010.

[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] Ferejohn, John A., Morris P. Fiorina, and Edward W. Packel. “Nonequilibrium solutions for legislative systems.” Behavioral Science 25.2 (1980): 140-148.

[17] Ferejohn, John A., Richard D. McKelvey, and Edward W. Packel. “Limiting distributions for continuous state Markov voting models.” Social Choice and Welfare 1.1 (1984): 45-67.

[18] Ferguson, John-Paul, and Sharique Hasan. “Specialization and career dynamics: Evidence from the Indian Administrative Service.” Administrative Science Quarterly 58.2 (2013): 233-256.

[19] 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.

[20] Fisman, Raymond, et al. ”Social Ties and the Selection of China’s Political Elite.” American Economic Review 110.6 (2020): 1752-81.

[21] Hall, Brian J., and Kevin J. Murphy. “Stock options for undiversified executives.” Journal of accounting and economics 33.1 (2002): 3-42.

[22] Hannes, Sharon. “Reverse Monitoring: On the Hidden Role of Employee Stock-Based Compensation.” Mich. L. Rev. 105 (2006): 1421.

[23] 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.

[24] Huang, Chien-Chung, and Telikepalli Kavitha. “Popular matchings in the stable marriage problem.” Information and Computation 222 (2013): 180-194.

[25] Huang, Chien-Chung, and Telikepalli Kavitha. “Popularity, mixed matchings, and self-duality.” Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms. Society for Industrial and Applied Mathematics, 2017.

[26] Iyer, Lakshmi, and Anandi Mani. ”Traveling agents: political change and bureaucratic turnover in India.” Review of Economics and Statistics 94.3 (2012): 723-739.

[27] Pérez-González, Francisco. “Inherited control and firm performance.” American Economic Review 96.5 (2006): 1559-1588.

[28] 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.

[29] Jenter, Dirk. “Understanding high-powered incentives.” (2001).

[30] Kavitha, Telikepalli. “A size-popularity tradeoff in the stable marriage problem.” SIAM Journal on Computing 43.1 (2014): 52-71.

[31] Kavitha, Telikepalli, and Meghana Nasre. “Optimal popular matchings.” Discrete Applied Mathematics 157.14 (2009): 3181-3186.

[32] Kavitha, Telikepalli, Julián Mestre, and Meghana Nasre. “Popular mixed matchings.” Theoretical Computer Science 412.24 (2011): 2679-2690.

[33] Kitchen, Martin. A Military History of Germany: From the Eighteenth Century to the Present Day. Bloomington: Indiana University Press, 1975.

[34] Koray, Semih. “Self-selective social choice functions verify Arrow and Gibbard-Satterthwaite theorems.” Econometrica 68.4 (2000): 981-996.

[35] Krehbiel, Keith. Information and legislative organization. University of Michigan Press, 1992.

[36] Lambert, Richard A. “Executive effort and selection of risky projects.” The RAND Journal of Economics (1986): 77-88.

[37] Lambert, Richard A., and David F. Larcker. “Stock options, restricted stock, and incentives.” Restricted Stock, and Incentives (April 2004) (2004).

[38] 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.

[39] Lazear, Edward P., and Paul Oyer. Personnel economics. No. w13480. National Bureau of economic research, 2007.

[40] Lazear, Edward P., and Sherwin Rosen. ”Rank-order tournaments as optimum labor contracts.” Journal of political Economy 89.5 (1981): 841-864.

[41] Lemos, Renata, and Daniela Scur. “All in the family? CEO choice and firm organization.” (2018).

[42] Longenecker, Clinton O., and Dennis A. Gioia. ”Neglected at the top-executives talk about executive appraisal.” MIT Sloan Management Review 29.2 (1988): 41.

[43] Luthans, Fred, Richard M. Hodgetts, and Stuart A. Rosenkrantz. Real managers. Ballinger, 1988.

[44] Maggi, Giovanni, and Massimo Morelli. “Self-enforcing voting in international organizations.” American Economic Review 96.4 (2006): 1137-1158.

[45] Manlove, David. Popular Matchings, Chapter 7 in Algorithmics of matching under preferences. Vol. 2. World Scientific, 2013.

[46] Markovits, Daniel. The meritocracy trap: How America's foundational myth feeds inequality, dismantles the middle class, and devours the elite. Penguin Books, 2020.

[47] 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.

[48] Mayhew, David R. Congress: The electoral connection. Yale university press, 1974.

[49] McCutchen, Richard Matthew. "The least-unpopularity-factor and least-unpopularity-margin criteria for matching problems with one-sided preferences." Latin American Symposium on Theoretical Informatics. Springer, Berlin, Heidelberg, 2008.

[50] McDermid, Eric, and Robert W. Irving. "Popular matchings: structure and algorithms." Journal of combinatorial optimization 22.3 (2011): 339-358.

[51] McKelvey, Richard D. "Intransitivities in multidimensional voting models and some implications for agenda control." (1976).

[52] Messner, Matthias, and Mattias K. Polborn. "Voting on majority rules." The Review of Economic Studies 71.1 (2004): 115-132.

[53] Meulbroek, Lisa K. "The efficiency of equity-linked compensation: Understanding the full cost of awarding executive stock options." Financial management (2001): 5-44.

[54] 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.

[55] Pérez-González, Francisco. "Inherited control and firm performance." American Economic Review 96.5 (2006): 1559-1588.

[56] Podolny, Joel M., and James N. Baron. "Resources and relationships: Social networks and mobility in the workplace." American sociological review (1997): 673-693.

[57] 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.

[58] Shubik, Martin. "Games of status." Behavioral Science 16.2 (1971): 117-129.

[59] Smith, Ryan A. "Race, gender, and authority in the workplace: Theory and research." Annual Review of Sociology 28.1 (2002): 509-542.

[60] 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).

[61] Sng, Colin TS, and David F. Manlove. "Popular matchings in the weighted capacitated house allocation problem." Journal of Discrete Algorithms 8.2 (2010): 102-116.

[62] Sonmez, Tayfun, and Tobias B. Switzer. "Matching With (Branch-of-Choice) Contracts at the United States Military Academy." *Econometrica* 81.2 (2013): 451-488.

[63] Thakur, Ashutosh. "Combining social choice and matching theory to understand institutional stability." (2020).

[64] Thakur, Ashutosh. "Matching Politicians to Committees." (2017).

[65] Thakur, Ashutosh. "Matching in the Civil Service: A Market Design Approach to Public Administration and Development" (2019)

[66] Villalonga, B. and Amit, R. (2006), 'How do family ownership, control and management affect firm value?', Journal of Financial Economics 80, 385–417

[67] Weitzman, Martin. "The Share Economy: Conquering Stagflation." Cambridge, MA: Harvard University Press (1984)

Appendix A. **Generalized voting rules: decisive coalitions.**

In Section 8, we characterized the effects of numerous commonly used voting rules. Here, we provide results characterized in terms of the sets of decisive coalitions associated with the aggregation rule.

A coalition of agents $D \subseteq (1, ..., N)$ is decisive if for any two allocations A and A', if $\forall i \in D$, $U_i(A') > U_i(A)$ implies the collective choice is for A' over A. For any aggregation rule 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 aggregation rule 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. Austen-Smith and Banks (2000)'s Lemma 2.2 (p. 42) shows that for an aggregation rule, the set of decisive coalitions is monotonic and proper.

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

In this section we provide results showing the robustness of Theorem 1 to aggregation rules that are *not collegial*.

Theorem 3. *If $\alpha + \beta = 1$, $\alpha, \beta > 0$, the alternative allocation protocol is restricted to swaps, and no agent belongs to every decisive coalition, i) the unique stable allocation is the first-best and ii) the first-best is dynamically reached with probability 1.*

Proof: i) Since no agent belongs to every decisive coalition, the meritocratic first-best allocation is stable because any swap reduces output, thereby making all $N-2$ uninvolved agents and at least one of the swapped agents worse off. There is no other stable allocation because every other allocation has at least one inverted 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. ii) Any swap that reduces output would make all $N-2$ uninvolved parties and at least one of the swapped agents oppose the proposed swap; if output remained constant then all $N-2$ uninvolved parties would be indifferent and one of the swapped agents would oppose. Hence, since no single agent belongs to every decisive coalition, the only swaps that pass are output-improving changes. Since the set of possible matchings is finite, this implies the first-best is reached with probability 1.

Unanimity rule and dictatorship are examples of voting rules where an agent belongs to every decisive coalition. Of course, we can always state weaker results that have $Q \leq N$ agents belonging to every decisive coalition: then the first best is achieved if and only if these Q agents are the highest quality and the initial allocation has these Q agents correctly occupying their highest status seats.

In fact, our main result is robust to even more complex voting rules in which voting power is based on which agent is assigned to which seat, i.e., allowing decisive coalitions to depend on the matching.

Theorem 4. *If $\alpha+\beta = 1$, $\alpha, \beta > 0$, the alternative allocation protocol is restricted to swaps, and no incorrectly matched (relative to his first-best assignment) agent-seat pair belongs to every relevant decisive coalition*[22]*, i) the unique stable allocation is the first-best and ii) the first-best is dynamically reached with probability 1.*

> **Proof**: 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 output-improving swaps and votes against output-decreasing swaps. Technical note i) 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 ii) 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.

These theorems show that though our main result, Theorem 1, is based on a stylized representation of voting, its conclusion — the convergence to and stability of the first-best — is robust to alternative specifications of influence, such as informal debate and deliberation, which arise organically. The conclusion's robustness and the model's generality is especially apparent in Theorem 3, where how much weight an agent is able to impose in the reorganization process can depend on the seat s/he currently occupies. (For example, an agent's influence over personnel reorganization can depend on whether he/she is a worker, manager, or CEO.)

[22]There are two technical notes: i) Relative to the status quo allocation, infeasible decisive coalitions are dropped when we consider the set of "relevant decisive coalitions." Namely, 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. ii) From each possible allocation, the set of relevant decisive coalitions should be nonempty.

Appendix B. Towards More Sophisticated Voting

Thus far we have assumed myopic voting. However, the strong organizational drift towards more meritocratic allocations and higher output (provided agents do not put zero weight on output) does not hinge on that assumption. In this section we show that Theorem 1 is robust to a variety of more sophisticated strategic voting behaviors.

B.1. **Forward-looking considerations for possible swaps 1-period ahead.**

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

When considering swapping agents j and k this period, the set of possible swaps in next period is, for $\beta > 0$, affected only for those agents whose quality lies between the j and k and who occupy seats in the status hierarchy which are between or below the seats where j and k are being swapped. 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-optionality based on forward-looking behavior disciplines the organization from correcting 'too quickly' and in 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. However, 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 set B to vote against. Hence such forward-looking behavior can only slow down convergence to the first best under $\beta > 0$.

B.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 $\alpha, \beta > 0$. An agent i votes for a proposed output-increasing change 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.

Agents are not fully forward-looking under short-/far-sighted voting voting: the organization might end up with an allocation that is stable because it is preferred by a majority to the first best. In such cases, the organization gets stuck, but this is unanticipated from a voting point that imposes a short-run requirement (tomorrow must be better than today) as well as a long-run criterion (the end-state must be better than today). Given this complex combination of short- and long-run considerations, the stable allocation that is reached depends on the initial allocation, the agenda of alternative allocations proposed, and the relative weights placed on α and β.

Let us investigate the effect of the weight on organizational output β on the stability of different allocations. We denote by A_s, the subset of all allocations that is stable and let $A_s(\beta)$ denote the stable subset of allocations given a particular value of β.

Proposition 8. *Assume short-/far-sighted voting. If $\alpha+\beta=1$, $\alpha,\beta>0$ and the alternative allocation protocol is restricted to swaps, then i) there exists $\overline{\beta}<1$ such that for all $\beta>\overline{\beta}$, the only stable allocation is the first best and it is reached with probability 1, ii) there exists a $\underline{\beta}>0$, such that for all $\beta<\underline{\beta}$, there exist stable allocations that are suboptimal, and iii) if $\beta_2>\beta_1>0$, then $A_s(\beta_2)\subseteq A_s(\beta_1)$.*

Proof: 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 long-run payoff criterion is always satisfied and any output-improving swaps will be passed. ii) Set $\beta=0$. If the status quo allocation is $(2,...,N,1)$, 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. iii) 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 the status, so again the shortfall in organizational output must be more than compensated for by the private payoff in status. Hence if an allocation is stable given β_2 then it must also be stable given β_1, thereby $A_s(\beta_2)\subseteq A_s(\beta_1)$.

B.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 represent.

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 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 $Y(A) > Y(B)$ 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:

- *Proof of ii)*:
  When $\beta = 0$, all possible assignments lead to an equivalence with respect to relabeling agents, as only status matters ($\alpha = 1$) and the environment is completely rivalrous given 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. Hence the 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, and hence belongs to the uncovered set. For $\beta = 0$, the proof of ii) above implies first best is also uncovered.

Hence, consider the remaining case of $\beta \in (0,1)$. Label the first best allocation X. To show that X is uncovered, we will construct an allocation Z such that for any allocation Y, $Z \succ Y$, but $X \succ Z$.

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, construct allocation Z as follows: 1) Leave all unchanged agents where they are in Y (note: they will prefer Z over Y as output increases), 2) Take the $\frac{N-1}{2}$ worst quality displacers, and move them back to their ideal positions (note: they will be made worse off compared to Y), 3) Keep the highest quality excess displacers where they are (note: they will prefer Z over Y as output increases), and 4) Order all displaced meritocratically in the seats that remain. (note: there are two cases, either: meritoratic ordering of displaced keeps them in the same position they were in with Y and they prefer Z as output increases OR meritocratic ordering of displaced moves them up in which they prefer Z as they get more output and higher status). Hence we have constructed $Z \succ Y$.

Now for the comparison of Z and X: First note that output goes down in Z by definition as X is first best. 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. And (2) By definition, 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. (1) and (2) alone make a minimal majority, hence $X \succ Z$.

- *Proof of iii)*:

  When $\beta = 1$ all agents strictly prefer the any allocation A with higher output than any allocation B. Moreover, the first best is unanimously preferred to any other allocation. Since this is a strict preference, by continuity, it continues to hold for β sufficiently close to 1.
- *Proof of iv)*:

  This holds trivially for $\beta = 0$ from ii). Hence consider $\beta > 0$.

  We want to show that the worst output, most un-meritocratic allocation we label 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 the worst output allocation 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 happen to be the relatively worst quality agents as Y is the worst output assignment) in the same seats, 3) Move the remaining lowest-seeded displacers in Y (who happen to be the relatively best relative 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, so might vote against Z for strictly positive β.

Moreover, we also get $X \succ Z$. The top $\frac{N-1}{2}$ seeded displacers can be better off as they improve status in Z compared to 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. 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 unmeritocratic moves comparing X to Z. Therefore, output under X is less than output under Z. Hence, while 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.

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.[23] However, simulations suggest that counterexamples are rare.[24]

[23]Consider the following counterexample: $N = 5$, agents' qualities are parameterized by 5,4,3,2,1 where a higher quality is better. Seat status of the 5 seats are parameterized by 5,4,3,2,1. $Y(A) = \sum_{i \in N} (SeatStatus_i)^2 * Quality_i$, $\beta = 0.7$, and all agents get utilities $[50, 40, 30, 20, 10]$ from being assigned to the seats respectively. Notation: we write allocation to be what status seat is given to each person, ordered by decreasing quality (e.g., 45321 means best quality agent is given seat of status 4, second-best agent gets the highest status seat 5, etc.). 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 respectively.

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

Appendix C. Proof of Proposition 1 from Thakur (2020)

This appendix reproduces, verbatim, the proof of Proposition 1 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 2 in our paper. Since Proposition 1 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,...,N to be assigned to N agents $1, ..., N$. Suppose each agent's preference over seats is identical, say $1 \succ 2 \succ ... \succ N$. Fixing the order of seats[25] to be $1, ..., N$. An allocation hence defines a permutation of agents over this ordered sequence of seats. For example, an allocation denoted by $5, 2, 9$,... implies that agents $5, 2, 9$,... were assigned seats 1,2,3,... 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 is 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 $1, 2, ..., N$ to $N, 1, 2, ..., N-1$. A $1:2$ transposition is the movement from $1, 2, ..., N-1, N$ to $2, 1, 3, ..., N$. That the set of all possible permutations, S_N, is generated (i.e., a successive composition of these two operators) using these two movements, is a theorem in abstract algebra for finite symmetric groups, stated usually as $(1, 2)(1, 2, ..., N)$ 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 1[26]. (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).

[25]This 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).

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

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

Fixing the sequence of seats $1, ..., N$, the figures show the sequences of majority-approved allocations $A_0, A_1, ...$ 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 $A_1, ..., A_N$, 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 $A_1, ..., A_{\frac{N}{2}}$, the bottom $N-2$ agents are moved to the top, hence approved by $N-2$ agents.

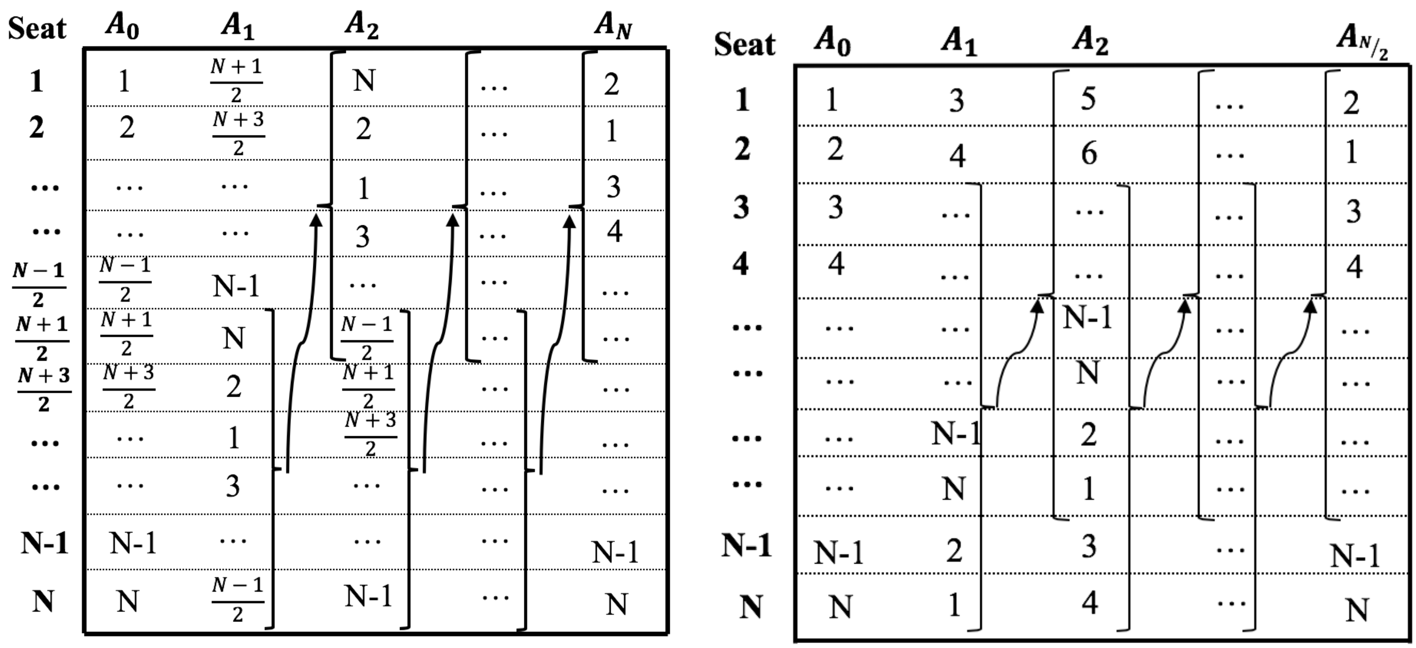

”