

Meritocracy and Homophily in Collegial Organizations

Paul-Henri Moisson^{†,*}

Jean Tirole^{†,**}

January 18, 2026

Abstract

In collegial organizations, homophily jeopardizes the recruitment of the most talented (“meritocracy”). We analyze the dynamics and welfare properties of an organization whose incumbent members coopt new ones in a forward-looking manner. We identify organizations in which meritocracy is likely to give way to favoritism and entrenchment, and investigate policy interventions (such as affirmative action, competitive rewards, discretionary overruling of majority decisions, blind or semi-blind hiring processes) and their unintended consequences.

Keywords: Cooptation, collegial organizations, Markov games, meritocracy, homophily, affirmative action, glass ceiling.

JEL Codes: D7, C73, D02, M5.

[†]The authors are grateful to Daron Acemoglu, Guido Friebel, Bob Gibbons, Bård Harstad, Johannes Hörner, Alessandro Pavan, Martin Peitz, Jérôme Renault, Patrick Rey, three anonymous referees, and to participants at numerous seminars and conferences for helpful suggestions. Funding from the European Research Council (grant agreement no. 101098319 - ERC HWS) is gratefully acknowledged.

^{*}Paris School of Economics and Ecole des Ponts ParisTech, paul-henri.moisson@psemail.eu.

^{**}Toulouse School of Economics (TSE), jean.tirole@tse-fr.eu.

1 Introduction

The selection of new members of a board of directors, a corporation, a cooperative, a trade or monetary union, an academic department or a polity, underlies institutional dynamics and determines whether the organization succeeds or is consigned to oblivion. Most often, new members are coopted,¹ i.e., recruited by incumbent members according to a (formal or informal) voting process, perhaps under constraints imposed by internal rules or external intervention. Even in more hierarchical organizations, recruitment often reflects some degree of collegiality, whereby whoever holds the formal authority on the recruitment decision puts substantial weight on subordinates' (candidates' would-be colleagues) opinions. We aim at studying the consequences of such collegiality.

This paper analyzes the Markovian dynamics, the discrimination in hiring (and promotion), and the welfare properties of an organization whose members are forward-looking and are driven by two motives in their cooptation decisions: talent and homophily. All else being equal, all members prefer a more talented candidate to a less talented one. However, homophily along a "horizontal" trait (gender, religion, ethnicity, politics, scientific field or approach, values, family, friendship, class loyalty...) makes members prefer an in-group candidate to an out-group one unless the latter is substantially more talented than the former. This misalignment of horizontal preferences creates a benefit from controlling the organization's recruitments. Members of the majority group may favor an in-group candidate over a more talented out-group one, so as to stuff the organization with their chums and thereby entrench their grip on the organization's recruitments. This violates meritocracy, defined here as the selection of the more talented.

The first contribution of our paper is to derive new theoretical predictions about the drivers of entrenchment in the absence of public intervention. *(a) Homophilic intensity.* An obvious factor is the preference for homophily: agents who attach more importance to mingling with their in-group are more likely to give up on meritocracy. *(b) Majority size and discrimination.* With linear or concave homophily benefits, thinner majorities discriminate more than larger ones. By contrast, with strictly convex homophily bene-

¹We focus on "cooptation" in the sense of "periodic selection by incumbent members of new members, according to a given voting rule". A second and equally important acception of "cooptation", associated with Selznick (1948, 1949), argues that absorbing new elements in an organization can be a means of averting threats to its stability or existence. We refer to the literature building on Acemoglu and Robinson (2000)'s celebrated analysis on the extension of the franchise to avoid upheaval (threat-averting cooptation involves the entire threatening group in Acemoglu-Robinson, and only a sub-group in Bertocchi-Spagat 2001).

fits, larger majorities may discriminate more than thinner ones. *(c) Uncertainty about future control.* A higher uncertainty for the majority on its ability to retain control, due for instance to uncertain voting participation or uncertain group allegiance, can lead to a more intense discrimination against minority candidates as the majority builds a precautionary buffer against unexpected losses of control. This "buffer effect" drives some unintended consequences of policy interventions (see below). *(d) Patience.* When facing a trade-off between coopting a minority member or a less talented majority one, the majority members weigh the discounted benefits from the talent differential (in the current recruitment) against those from future control (affecting all future hires). The latter increase faster than the former as the discount factor increases: a longer expected tenure (a longer time-horizon) within the organization makes majority members more prone to entrench. *(e) Weak-link principle.* It takes two to build meritocracy: the cooptation process is only as meritocratic as the group with the strongest homophily preferences, or lowest patience allows it. *(f) Homogamic evaluation capability.* Talent evaluation that is more accurate within the in-group (due to field expertise or familiarity) under weak conditions raises entrenchment, as the majority is less able to identify talent within the pool of minority candidates.

As allowed by our framework, entrenchment is not always socially detrimental – for instance, friendship circles are often based on homophily in tastes. We emphasize a fundamental asymmetry between entrenchment and meritocracy, though: When equilibrium behavior leads to meritocracy, meritocracy does dominate entrenchment in terms of welfare. By contrast, when equilibrium behavior leads to entrenchment, meritocracy may yield a higher welfare than entrenchment. We interpret this discrepancy as a *collegial bias against meritocracy*.

Our second contribution is an analysis (the first to the best of our knowledge) of familiar public policy interventions to promote meritocracy.² Policy interventions are motivated by the misalignment between majority and social interests when the majority is tight. The majority trades off a higher quality brought by a more talented minority's hire against a probabilistic, but repeated loss of future homophily benefits. Two distinct families of interventions aim at altering this tradeoff in the direction of meritocracy:

(i) Interventions reducing the value of control rights. The first family of interventions makes it less appealing for the majority to secure control over future decisions. A

²The different families encompass different information structures and means of interventions.

diminished value of control rights down the road tilts the balance in favor of a meritocratic choice, at least if it does not deteriorate the minority’s payoff. We examine three such policies: equal treatment of members, affirmative action (minority quotas), and homophily-blind selection. In each case, the intervention favors meritocratic choices by a tight majority. However, by not allowing the majority to exercise free, informed choice, these policies unfortunately come with side effects: inefficient internal policies (equal treatment of members), reverse discrimination (quotas), and a loss of quality information (blind selection).

(ii) Interventions tilting the current decision toward meritocracy. The second family of interventions makes today’s meritocratic choice more self-evident to the majority. It is composed of two subgroups of interventions: command-and-control policies, and incentive schemes. Command-and-control policies include semi-blind recruitment (introducing an initial homophily-blind stage in the selection process for candidates, thus implying that the short list selected at the first stage for stage-2 interviews by the majority may be limited to minority candidates), as well as the previous fully blind procedure,³ and micro-management (the principal sporadically imposes the choice of a more talented minority candidate).

The second subgroup of interventions tilting the current decision toward meritocracy is based on incentives. We cover three such policies, associated with either financial or symbolic rewards. All are zero-sum/budget neutral: Financial rewards for quality (whether provided through research assessment exercises, research councils’ grants or market-based rewards) are financed through a lower base-resource level; similarly, fines for discriminating against minority candidates go into a general increase of resources for the organization; reputational (symbolic) rewards are boosted by making it more salient that discrimination is happening.

One might expect command-and-control and incentives to be equivalent, as they tilt the current decision in the same direction. But in a dynamic environment with forward-looking players, they induce different choices before the majority is at risk (“up the road”). The command-and-control approach increases the majority’s uncertainty about future control, thereby increasing its willingness to avoid the dangerous zone of tight majorities. It encourages the formation of a precautionary buffer (which, except for homophily-blind hiring, the majority has the means to build). Such super-entrenchment

³Homophily-blind selection does two things simultaneously: (a) it annihilates the value of control rights, but (b) it also makes discrimination impossible in the first place as the majority does not observe the horizontal types. So, this intervention can be classified both as reducing the value of control rights, and tilting the current decision toward meritocracy (command-and-control).

induces an important leakage in the policy: by inducing more meritocracy today, the policymaker increases discrimination earlier on.⁴ By contrast, the “Pigovian” incentive-based policies that align the majority’s interests with society’s do not induce the buffer effect associated with command-and-control, nor the decision-inefficiency side effects associated with interventions reducing the value of control rights.

Technical contribution and roadmap.

Section 2 builds the baseline model under laissez-faire. As we later show, this model gives the best chance to meritocracy by assuming that (i) the majority can perfectly predict prospective hires’ allegiance and there is no uncertainty about turnout in future recruitment elections, and (ii) groups can identify talent equally well for out-group candidates and for in-group ones.

The organization has an arbitrary size. There are two horizontal groups, and two talent levels (we later generalize the talent distribution to a continuum). Organization’s members enjoy linear quality and homophily benefits from their colleagues’ attributes (we later generalize to non-linear and asymmetric homophily benefits). Quality benefits exceed homophily ones (the interesting case, as otherwise the majority only hires in-group candidates). Section 2 fully characterizes pure-strategy Markov Perfect Equilibria (MPEs). We show that MPEs satisfy the following properties: (a) equilibrium strategies are *meritocratic*, except perhaps for tight majorities (when a minority appointment may lead to a loss of control), in which case *basic entrenchment* may arise; (b) a group is more inclined to be meritocratic if the other group also is (strategic complementarity); (c) in the symmetric case, the organization is either meritocratic or basically entrenched, and the two regimes coexist over a non-empty range of the quality-over-homophily-benefit ratio.

Section 2 then computes the welfares of current minority and majority members. The two equilibria, when coexisting, are Pareto-ranked with meritocracy dominating entrenchment, which enables us to make a selection and perform comparative statics and policy evaluation. A second measure of welfare is aggregate ergodic welfare, which, by focusing on long-term payoffs, embodies future recruits’ rents on top of current members’s value functions. With this criterion as well, the meritocratic equilibrium, while delivering lower homophily benefits on average, dominates the entrenchment one.

⁴In the case of blind hiring, the majority does not have the ability (the information) to super-entrench (or even entrench). The inefficiency then stems from the loss of vertical information induced by the “blinding”.

Section 3 studies the (aforementioned) policy interventions. Section 4 relaxes the assumptions of the baseline model, emphasizing the robustness of the analysis in Sections 2 and 3. It considers asymmetric preferences (either for homophily or time horizon), and establishes a "weak link" principle for meritocracy (Section 4.1). It extends the analysis to a continuous quality space (4.2), nonlinear homophily benefits (4.3), homogamic evaluation capability (4.4.1), uncertain voting participation (4.4.2), and "anterooms for appointments", which can be external when rejected candidates may reapply or internal when junior members may be promoted to a senior position⁵ (4.5). The paper concludes by discussing the related literature (Section 5) and avenues for future theoretical and empirical research (Section 6). Omitted proofs can be found in the Online Appendix.

2 Baseline model

There is an infinite time horizon with periods $t \in (-\infty, +\infty)$. The organization is composed of $N = 2k$ members. At the beginning of each period, one member of the organization, drawn randomly from the uniform distribution, departs. We denote by δ the "life-adjusted discount factor", i.e. the pure-time discount factor times the probability of still being a member of the organization in the following period: letting $\delta_0 \in (0, 1)$ denote the pure-time discount factor, then $\delta \equiv \delta_0(1 - 1/N)$. The departure is immediately followed by a recruitment. The intra-period timing is summarized in Figure 1.

Each individual has a two-dimensional type. The vertical type captures talent or quality and takes one of two possible values, 0 (mediocre) or \tilde{s} (talented), where $\tilde{s} > 0$ is the incremental per-period contribution of a talented individual to each other member's payoff. The horizontal type stands for race/gender/tastes/opinions and can take two values $\{A, B\}$. A member of horizontal type $X \in \{A, B\}$ exerts per-period externality $\tilde{b}_X > 0$ on members of the same type,⁶ but not on members of the opposite type, and this regardless of their talent.⁷

⁵In hierarchical organizations, the oft-made observation that minorities experience difficulties in rising above a certain level suggests that meritocracy is more often violated at higher than at lower levels. Even if in-group favoritism contributes to discrimination against minorities, it is not a priori obvious that it should imply a lower rate of promotion for the latter (a "glass ceiling"). Nonetheless, Section 4.5.2 shows in the natural extension to a two-level organization that a glass ceiling results from control being located at the senior level.

⁶The case $\tilde{b}_X < 0$, corresponding to *negative homophily* – e.g., envy towards the likes, preference for diversity or for a smaller in-group, etc. (see for instance Bagues and Esteve-Volart 2010) – can be accommodated in our model. See Online Appendix S.

⁷Members may enjoy direct homophily benefits, associated with the desire of sharing identity (political or other) or interests (say, similar leisure activities) with fellow members. Alternatively, homophily benefits may be more instrumental/indirect. Having like-minded members on board allows one to weigh

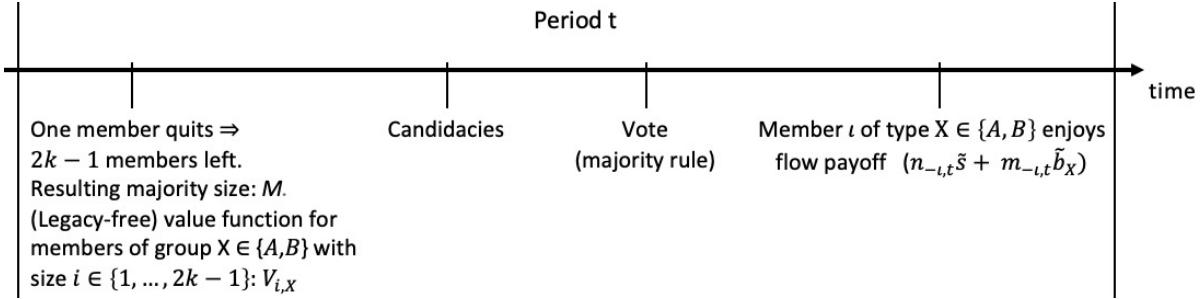


Figure 1: Timing.

We thus assume that each member derives utility from: (i) their colleagues' talent, i.e. the vertical attributes of members of the organization, and (ii) homophily over tastes: *ceteris paribus*, each member prefers colleagues who share their horizontal type. So, member i of group $X \in \{A, B\}$ receives date- t flow payoff

$$u_{i,t} = n_{-i,t} \tilde{s} + m_{-i,t} \tilde{b}_X$$

where $n_{-i,t} \leq N - 1$ and $m_{-i,t} \leq N - 1$ are respectively the number of talented colleagues and the number of in-group colleagues of member i , at date t .⁸

In each period, there is at least one candidate of each type. Assuming then that there is exactly one of each type involves *no loss of generality* as all members of the organization prefer the most talented candidate of a given horizontal type to any candidate of the same type but with lesser talent, and are indifferent if there are multiple "most-talented" candidates of a given horizontal type. Moreover, we assume that candidates apply to become members only once and that the candidates' types are observable prior to the vote – we will later relax these two assumptions.

Let $s \equiv \tilde{s}/[1 - \delta_0(1 - 2/N)]$ denote the expected incremental lifetime contribution of a new talented (relative to mediocre) appointee to each incumbent member of the organization.⁹ We similarly denote by $b_X \equiv \tilde{b}_X/[1 - \delta_0(1 - 2/N)]$ the expected lifetime homophily benefit for an incumbent member of group $X \in \{A, B\}$ generated by a new in-group member.

The decision rule is the majority rule, with each of the $2k - 1$ members of the organization at the time of the vote having one vote. Let $M \in \{k, k + 1, \dots, 2k - 1\}$ denote the size of (number of individuals in) the majority. We say that the majority is *tight* if

⁸on organizational decisions and the sharing of private benefits: more committees are filled by in-group members and more suggestions favorable to the group are made.

⁹Our insights are unchanged if a talented member derives a "quality payoff" from their own talent.

⁹The term $\delta_0(1 - 2/N)$ stems from the conditioning on both the current member and the newly recruited one still being in the organization in the next period.

$M = k$.¹⁰

To make things interesting, we assume $s > b_X$ for all $X \in \{A, B\}$. Otherwise, if $b_X \geq s$, systematically voting for the majority candidate would yield the highest possible continuation payoff for majority X , and (in the absence of coordination failure) such a majority would always move towards perfect homogeneity.

Candidates' talents are i.i.d. across periods. We let $x \in (0, 1/2]$ denote the probability that the majority (or minority) candidate is more talented (i.e. has vertical type s while the other candidate has vertical type 0), and thus $(1 - 2x)$ is the probability that they are equally talented (either both of quality s or both of quality 0). Let $\alpha \in [0, 1]$ denote the probability that both are of talent s conditional on both being equally talented. The probability of an in-group (or out-group) candidate being of quality s is thus equal to $\bar{x} \equiv x + (1 - 2x)\alpha$.

Our equilibrium concept is pure-strategy Markov Perfection. Given the other group's strategy, all members of a given group $X \in \{A, B\}$ at date t have the same objective function. Moreover, if they have not exited by date $t + \tau > t$, they will have the same date- $(t + \tau)$ current-plus-continuation payoff function as the other members of the same group, regardless of their respective cohort and talent. So, only the size M of the majority and its identity X are payoff-relevant in the sense of Maskin-Tirole (2001). Markov strategies therefore depend neither on time t , nor on the cohorts or talents of incumbent members.

In addition, we assume that each member votes as if they were pivotal, i.e., as if they alone chose the candidate. Hence, we ignore coordination failures in which, say, a majority member votes for an unfavored candidate because other majority members also do.¹¹ Consequently, at any date, all members of a given group vote unanimously to maximize their current-plus-continuation payoff.¹²

¹⁰We refer to a majority member as "he", to a minority member as "she", to a generic organization member as "they", and to the principal – whenever there is one – as "it".

¹¹The assumption that agents vote as if they were pivotal could stem in particular from a trembling-hand requirement as in Acemoglu et al. (2009), or from a coalition-proofness requirement among current members of the same horizontal group (majority/minority).

¹²Since we thus rule out coordination failures within the majority, the minority's current voting behaviour is for now irrelevant (uncertain voting participation or identification of group allegiance will be considered in Section 4.4.2).

2.1 Equilibrium characterization and existence results

2.1.1 Majority's best response and strategic complementarities

The state variable for a majority of horizontal type $X \in \{A, B\}$ is its size $M \in \{k, \dots, N-1\}$. We study the "best response" for a majority of type $X \in \{A, B\}$ and size $M \in \{k, \dots, N-1\}$ to the other group's strategy, summarized by the current majority's continuation value upon losing control (reaching size $k-1$).

Since the present discounted value of benefits accruing from other incumbent members plays no role in an MPE, we do not include the legacy terms in the expression of the value functions. For any group size $i \in \{1, \dots, N-1\}$ just before candidacies are declared (see Figure 1), we denote by $V_{i,X}$ the value function of an individual in group X : $V_{i,X}$ is the expected discounted value of flow payoffs brought about by colleagues who *will* be coopted later, in the current period and in future periods.

A majority member's continuation value at majority size $M \geq k$ is given by

$$b_X + s_{\text{maj}} + \delta \left[\frac{M}{N-1} V_{M,X} + \left(1 - \frac{M}{N-1}\right) V_{M+1,X} \right]$$

if the majority candidate with talent (expected lifetime contribution) $s_{\text{maj}} \in \{0, s\}$ is recruited in the current period, and by

$$s_{\text{min}} + \delta \left[\frac{M-1}{N-1} V_{M-1,X} + \left(1 - \frac{M-1}{N-1}\right) V_{M,X} \right]$$

if the minority candidate with talent (expected lifetime contribution) $s_{\text{min}} \in \{0, s\}$ is. The value function $V_{M,X}$ of a majority member at majority size M is the expectation of its continuation value over all current-period possible events (candidates' profiles and recruitment decisions).¹³ The majority's choice between the two candidates is thus determined by the following comparison:

$$b_X + s_{\text{maj}} - s_{\text{min}} + \delta \left[\frac{M-1}{N-1} (V_{M,X} - V_{M-1,X}) + \left(1 - \frac{M}{N-1}\right) (V_{M+1,X} - V_{M,X}) \right] \leq 0. \quad (1)$$

Definition (Canonical strategies).

(i) *The majority votes for the majority candidate if the latter is at least as talented as the minority candidate.*

¹³Moreover, any continuation value $V_{i,X}$ with $i \in \{1, \dots, N-1\}$ thus lies in the interval $[0, ((\bar{x} + x)s + (1 - x)b_X)/(1 - \delta)]$.

(ii) When the minority candidate is more talented, a type- X majority, with $X \in \{A, B\}$, votes for the majority candidate with probabilities $\{\sigma_X(M)\}_{M \in \{k, \dots, N-1\}}$ with either $\sigma_X(M) = 0$ for all $M \geq k$, or $\sigma_X(M) = 1$ for $k \leq M \leq k+l$ and $\sigma_X(M) = 0$ if $M \geq k+l+1$ for some $l \geq 0$.

We will say that a type- X majority is

- *meritocratic* if $\sigma_X(M) = 0$ for all $M \geq k$;
- *basically entrenched* if it favors a mediocre majority candidate over a talented minority one only when the majority is tight ($M = k$), i.e. if $\sigma_X(k) = 1$ and $\sigma_X(M) = 0$ for all $M \geq k+1$.
- *super-entrenched at level $l \geq 1$* if $\sigma_X(M) = 1$ for $M \in \{k, \dots, k+l\}$, and $\sigma_X(M) = 0$ for $M \geq k+l+1$,
- *fully entrenched* if it is super-entrenched at level $l = k-1$, i.e., if $\sigma_X(M) = 1$ for all $M \geq k$.

Online Appendix A proves the following intuitive property:

Lemma 1. (Majority's best response and canonical strategies) Fix $V_{k-1,X}$ in the feasible range ($[0, ((\bar{x}+x)s + (1-x)b_X)/(1-\delta)]$). The majority's best response to $V_{k-1,X}$ among pure Markov Perfect strategies is either meritocracy or basic entrenchment.

Intuitively, the assumptions of the basic model ensure that control can be retained simply by coopting a majority candidate when the majority is tight ($M = k$). Therefore, unless control is immediately at stake ($M = k$), the majority's strict best response is always to recruit the most talented candidate, breaking ties in favor of the in-group candidate.

Our next economic insight – the strategic complementarity of meritocracy and entrenchment – builds on Lemma 1. Let us, abusing notation, denote by $V_{i,X}^{r,r'}$ group X 's continuation value function when it has size i and follows strategy r when it has the majority, with $r = m$ if the strategy is meritocracy and $r = e$ if it is basic entrenchment, and the other group follows strategy $r' \in \{m, e\}$ when it has the majority. When $r = e$, control never switches and the other group's strategy is irrelevant, and so $V_{i,X}^{e,m} = V_{i,X}^{e,e}$ for all $i \geq k$.

When $r = m$ and the other group is basically entrenched, the expected payoff from the current hire in $V_{i,X}^{m,e}$ for $i \leq k-1$ is bounded above by $(\bar{x}+x)s + xb$ and is strictly

lower than the bound when $i = k - 1$, while the expected payoff from the current hire in $V_{i,X}^{m,m}$ for $i \leq k - 1$ is bounded below by $(\bar{x} + x)s + xb$ (and is strictly higher than this bound when $i \geq k$ if $x < 1/2$).¹⁴ This implies that $V_{k,X}^{m,m} > V_{k,X}^{m,e}$. And therefore,

$$V_{k,X}^{m,m} - V_{k,X}^{e,m} > V_{k,X}^{m,e} - V_{k,X}^{e,e}.$$

Consequently, group X is more inclined to be meritocratic when it has a tight majority if the other group is also meritocratic.

Proposition 1. (*Strategic complementarity*) *A given group $X \in \{A, B\}$ is more inclined to be meritocratic (resp. basically entrenched) if the other group is itself meritocratic (resp. basically entrenched).*

2.1.2 The symmetric case

Except in the asymmetric extension of Section 4.1 and for expositional conciseness, we henceforth restrict our attention to the symmetric case in which both horizontal groups have the same homophily preferences: $b_A = b_B \equiv b$.

In the symmetric case, Lemma 1 and Proposition 1 together imply that generically, any pure-strategy MPE is symmetric and either meritocracy or basic entrenchment,¹⁵ i.e., that the behaviors of A and B majorities are the same, either meritocratic or basically entrenched. To alleviate the notation in the symmetric case, we therefore drop the subscript X in the value function, and let $V_i^r \equiv V_i^{r,r}$ for all $i \leq k$ and $r \in \{m, e\}$.¹⁶ We thus refer equivalently to a meritocratic/basic-entrenched majority or organization.

Lemma 2 below states, and Figure 2 illustrates, the (monotonicity and increasing/decreasing differences) properties of the value functions of majority and minority members under the meritocratic and basic-entrenchment strategies.¹⁷

Lemma 2. (*Properties of value functions in the meritocratic (m) and basic-entrenchment (e) equilibria*)

¹⁴For $i \geq k$, the expected payoff from the current hire in $V_{i,X}^{m,e}$ is the same as the one in $V_{i,X}^{m,m}$, and the transition probabilities in the two continuation values are identical.

¹⁵To establish symmetry, suppose by contradiction that in equilibrium, group X is meritocratic and group Y is basically entrenched. Necessarily, $V_{k,X}^{m,e} \geq V_{k,X}^{e,e}$ and $V_{k,Y}^{e,m} \geq V_{k,Y}^{m,m}$. However, for any $r, r' \in \{e, m\}$, $V_{k,X}^{r,r'} = V_{k,Y}^{r,r'}$ as payoffs are symmetric, and we know from the proof of Proposition 1 that $V_{k,X}^{m,m} - V_{k,X}^{e,m} > V_{k,X}^{m,e} - V_{k,X}^{e,e}$. Hence, the equality $V_{k,X}^{e,e} = V_{k,Y}^{e,m}$ would imply that $V_{k,X}^{m,e} > V_{k,X}^{m,m}$, a contradiction.

¹⁶As we noted before, when the majority is entrenched, the other group's strategy is irrelevant, and thus, for $r' \in \{m, e\}$, $V_i^e = V_i^{e,r'}$ for $i \geq k$, while $V_i^e = V_i^{r',e}$ for $i \leq k - 1$.

¹⁷By "decreasing differences" (resp. "increasing differences"), we refer to the following concavity (resp. convexity) property: $|V_{i+1}^r - V_i^r| \leq |V_{j+1}^r - V_j^r|$ (resp. $|V_{i+1}^r - V_i^r| \geq |V_{j+1}^r - V_j^r|$) whenever $j < i$.

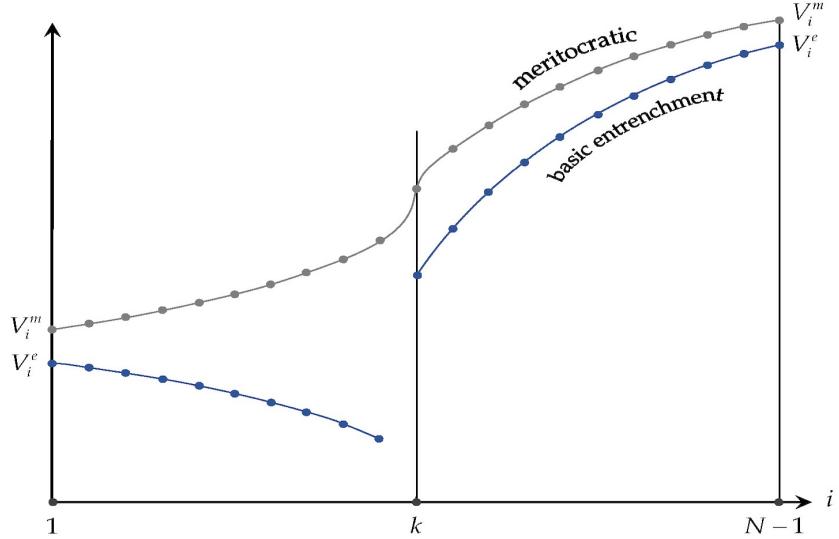


Figure 2: Properties of value functions in the (symmetric) meritocratic and basic-entrenchment equilibria.

- (i) (*Majority value function*) For $i \in \{k, \dots, N-1\}$, V_i^e is strictly increasing in i and has strictly decreasing differences. Similarly, V_i^m is increasing in i and has decreasing differences, strictly so if and only if $x < 1/2$.
- (ii) (*Minority value function*) For $i \in \{1, \dots, k-1\}$, V_i^e is strictly decreasing in i and has strictly increasing differences in i . By contrast, for $i \in \{1, \dots, k-1\}$, V_i^m is increasing in i and has increasing differences in i , strictly so if and only if $x < 1/2$.
- (iii) (*Control benefits*) For $r \in \{e, m\}$ and any $i \geq k$, $V_i^r \geq V_{N-1-i}^r$ (strictly so when $r = e$, and when $r = m$ and $x < 1/2$).

Intuitively, the three parts of Lemma 2 stem from the following observations. Firstly, in both the meritocratic and basic-entrenchment equilibria, the majority always picks its "myopically favorite" candidate except in the basic-entrenchment equilibrium when $M = k$, where "myopically favorite" refers to the choice the majority would make in the absence of future elections or, equivalently, if future hiring decisions did not hinge on the current one. The higher M is, the more remote the appointment of a myopically suboptimal candidate (basic-entrenchment equilibrium) and the more remote a possible loss of control (meritocratic equilibrium).

Secondly, for minority members, the impact of moving further away from the tight-majority state ($M = k$) depends on the equilibrium: in the basic-entrenchment equilibrium, the further away from minority size $k-1$, the smaller the additional *loss* of getting one step closer to the majority's entrenched recruitment at $k-1$, whereas in the merito-

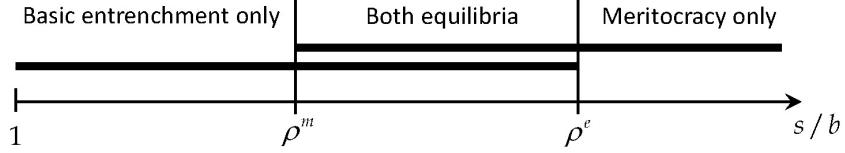


Figure 3: Existence regions for meritocratic and basic-entrenchment equilibria over the s/b line.

cratic equilibrium, the further away from minority size $k - 1$, the smaller the additional *benefit* of getting one step closer to possibly seizing control of the organization.

Thirdly, homophily induces a *benefit from control* for the majority whenever candidates have the same talent – as the majority can then pick its in-group candidate at no cost in terms of quality.¹⁸

Proposition 2. (*The symmetric case*)

- (i) All pure-strategy Markov Perfect equilibria are symmetric, and either meritocratic or basically entrenched.

There exist finite thresholds ρ^e and ρ^m satisfying: $1 \leq \rho^m < \rho^e < +\infty$, such that

- (ii) The basic-entrenchment equilibrium exists if and only if $s/b \leq \rho^e$,
- (iii) The meritocratic equilibrium exists if and only if $s/b \geq \rho^m$.
- (iv) Patience (in terms of pure-time preference) fosters entrenchment: for any δ_0 , $\partial \rho^m / \partial \delta_0 \geq 0$, and $\partial \rho^e / \partial \delta_0 \geq 0$.

Figure 3 describes the existence regions over the line s/b for given x, δ . For s/b close to 1, selecting an untalented peer over a better qualified minority candidate comes at little cost for the majority, and there is a benefit from keeping control, so the majority is basically entrenched. As the ratio quality/homophily benefits s/b increases, the cooptation game moves from a (bounded) region where only the basic-entrenchment equilibrium exists, to an intermediate (bounded) interval where, due to strategic complementarities, both equilibria coexist.¹⁹ As s/b continues to increase, it reaches the (half-line) region where only the meritocratic equilibrium exists.

¹⁸The benefit from control persists with a continuum of vertical types (see Section 4.2) as the majority then reaps a homophily benefit when recruiting its in-group candidate against a slightly more talented out-group candidate.

¹⁹As usual, there is then a third, mixed-strategy equilibrium with $\sigma(k) \in (0, 1)$. (Recall that we focus on pure-strategy equilibria throughout the paper. Besides, as will be clear shortly, this mixed-strategy equilibrium is dominated by the pure-strategy meritocratic equilibrium.)

As the discount factor δ_0 increases, the existence region of the meritocratic equilibrium shrinks while that of the basic-entrenchment equilibrium widens: The benefits from future control (affecting all future hires) increase faster than the benefits from the current recruit's talent, and so, when members become more patient, the relative cost of losing the majority increases.

Remark. If $x = 1/2$, i.e. the probability that both candidates have the same vertical type is nil, then $\rho^m = 1$: for any $s \geq b$, there exists a meritocratic equilibrium. The result is intuitive, as there is no benefit from control. By contrast, $\rho^m > 1$ whenever $x < 1/2$.

2.2 Welfare

2.2.1 Members' value functions: Meritocracy vs basic entrenchment

We first consider current members' welfare, defined as their expected discounted surplus (from current and future hires), at any given legacy and period, therefore computed from the continuation payoffs V_i^r .

The next Proposition shows that, when they coexist, the meritocratic equilibrium is preferred to the basic-entrenchment one by all current members of the organization. Intuitively, at any given majority size, minority members prefer the meritocratic equilibrium, while majority members, who can always select to be entrenched, weakly prefer the meritocratic equilibrium which delivers a higher payoff when surrendering control.²⁰

Proposition 3. (Members' value functions: Meritocracy vs basic entrenchment) *Whenever the meritocratic and the basic-entrenchment MPE coexist, i.e., for $s/b \in (\rho^m, \rho^e)$, at any majority size the meritocratic equilibrium is preferred by all current members of the organization to the basic-entrenchment equilibrium.*

As a consequence, from the perspective of current members, the meritocratic equilibrium Pareto-dominates the basic-entrenchment equilibrium whenever they coexist. This observation motivates our equilibrium selection – which we will use to perform our policy analysis (Section 3).

Assumption. (Equilibrium selection). *Whenever two equilibria coexist, coordination occurs on the meritocratic one. So, under laissez-faire, basic entrenchment prevails if and only if $1 < s/b < \rho^m$.*

²⁰Regardless to the regime $r \in \{m, e\}$, the majority faces an optimal stochastic control problem with boundary value V_{k-1} . All valuations V_{k+l} , with $l \geq 0$, are therefore non-decreasing functions of V_{k-1} .

2.2.2 Ergodic aggregate welfare

We now draw an aggregate-welfare comparison between basic entrenchment and meritocracy in their respective ergodic distribution from the perspective of a utilitarian principal or third-party putting at least as much weight on quality as on homophily benefits. We will indeed use this ergodic aggregate welfare to measure the impacts of different interventions in our policy analysis (Section 3): Denoting by S the organization's ergodic per-period aggregate quality, by B the ergodic per-period aggregate homophily benefits,²¹ the principal's objective writes

$$W \equiv qS + B$$

where $q \geq 1$ is the (relative) weight put by the principal on quality relative to homophily. The "no-externality-on-third-parties case" $q = 1$ corresponds to the maximization of (ergodic) total member surplus. But it often makes sense to assume that $q > 1$: homophily benefits are fully appropriated by the members, while talent yields benefits for both members and their organization or society (taxes, innovation, prestige, etc.).

Proposition 4. (*Ergodic per-period aggregate welfare*) *For any $s \geq b$, $W^m > W^e$, i.e., the meritocratic regime dominates the basic-entrenchment one in terms of ergodic per-period aggregate welfare.*

This result stems from the following intuition (see Online Appendix E for details): While basic entrenchment leads to larger majorities and thus larger homophily benefits, meritocracy delivers higher quality benefits, which not only have a higher social value "per unit" than homophily ones (as $qs \geq b$), but also accrue to *all* members whereas homophily ones profit only the in-group.

The higher aggregate welfare delivered by meritocracy motivates the policy interventions which we study in Section 3.

²¹Let ν_i^r denote the ergodic probability of majority size $i \in \{k, \dots, N\}$ at the end of a period in regime $r \in \{e, m\}$ (see Online Appendix E for its expression). Then, for regime $r \in \{e, m\}$,

$$B^r \equiv \sum_{i=k}^N \nu_i^r \left[i(i-1) + (N-i)(N-i-1) \right] \tilde{b},$$

while

$$S^m \equiv N(N-1)(\bar{x} + x)\tilde{s} \quad \text{and} \quad S^e \equiv N(N-1) \left[\nu_{k+1}^e \frac{k+1}{N} \bar{x} + \left(1 - \nu_{k+1}^e \frac{k+1}{N} \right) (\bar{x} + x) \right] \tilde{s}.$$

Hence, in particular, $S^m - S^e = N(N-1) \left(1 - \nu_{k+1}^e \frac{k+1}{N} \right) x\tilde{s}$.

The collegial bias against meritocracy. Propositions 2 and 4 point at an important asymmetry between entrenchment and meritocracy. While for any $s/b \geq 1$, meritocracy dominates basic entrenchment in terms of (ergodic aggregate) welfare, equilibrium behavior leads to basic entrenchment for any $s/b \in (1, \rho^m)$. By contrast, for $s/b < 1$, equilibrium behavior leads to full entrenchment, which does dominate meritocracy in terms of welfare. We interpret this discrepancy as a *collegial bias against meritocracy*.

2.3 Empirical evidence

There is growing evidence about the failure of meritocracy in various organizational environments.²² An important prediction of our analysis (and of our later extensions to super-entrenchment) is that weaker majorities discriminate more than stronger majorities. This is consistent with a number of empirical studies. Indeed, demographic change – and in particular increased local out-group presence – often shifts political behavior, voting and policy outcomes (Hopkins 2010, Dustmann et al 2019). More specifically, Behrens et al (2003) and Reny-Newman (2018), studying local demographic changes, show that rising minority sizes lead majorities to vote for more majority-preserving propositions (such as felon disenfranchisement when the proportion of felons is higher in the minority, and housing discrimination). Conversely, Enos (2016) shows that lower minority sizes lead to less conservative voting from majority voters. Similarly, Bagues et al. (2017) find that in (Italian and Spanish) scientific committees, male evaluators become less favorable to women if a woman joins the evaluation committee, suggesting horizontal control concerns from male evaluators.²³

Importantly, our results rely only on control considerations, and not on increased prejudice against minorities: We are not aware of empirical work that directly compares the sensitivity of behavior and prejudice to out-group size. Combined with evidence that contact can actually lower prejudice (Enos 2014), this pattern suggests that observed behavioral responses stems from a control logic, as in our model, rather than from rising animus.²⁴

²²For instance, Rivera (2012) finds evidence of biased hiring based on shared leisure activities, while Zinovyeva-Bagues (2015) shows that in the Spanish centralized process for promoting researchers to the ranks of full and associate professor, the promotion rate is higher when evaluated by the PhD advisor, a colleague or coauthor, and that the evaluation bias dominates the informational gain.

²³This evidence is also consistent with our study of uncertain participation/uncertain group allegiance (Section 4.4.2), which emphasizes that more insecure majorities are more prone to (super-)entrench.

²⁴In fact, prejudice may even reflect motivated beliefs and mirror control considerations: To reduce feelings of guilt from discriminating against a more talented minority (for, people derive utility from believing they do the morally right thing), majority members may persuade themselves that their discriminating against minorities only reflects the fact that the latter *deserve* their treatment (see the

3 Policy

3.1 Policy instruments and principal's objective function

Our policy analysis encompasses the most common real-world approaches. Before studying their consequences, we briefly describe the information and means they require to be implemented, and so the different environments in which they arise.

Observable variables. Policy instruments hinge on observability. The principal – e.g., a government, a regulating authority, an NGO able to engineer social pressure, or a major shareholder able to influence the organization's internal rules – may observe *talent*, at least sporadically and usually at a cost: the date- t relative quality of current candidates, or else the overall quality S_t of the organization. Assessments of S_t may be provided either within a research excellence framework evaluating the impact of research institutions (such as the ones operative in the UK, Germany and France) in academia, or by a market mechanism bringing money and prestige to a successful institution in the corporate and NGO world.

The principal may alternatively observe the degree of *homogamy*, again either of the current hire with the majority, or of the overall homogamy, M_t , of the organization. In some cases, the observation of M_t may be costly, but straightforward; gender homogamy is a case in point. In other cases, it may be difficult, even impossible for an outsider to access a reliable and verifiable measure of homogamy: homophilic tastes may result from common leisures, culture, religion, old-boy networks, field preferences and so on. We distinguish three broad groups of interventions depending on the principal's information and means: micromanagement, macromanagement, and internal governance (see Figure 4).

(a) *Micromanagement* refers to interventions in hiring decisions on the basis of information about the relative talent of the current candidates, the fact that the selected candidate belongs to the majority group (homogamy information), or even no information at all. Micromanagement interventions include discretionary overrulings (the principal can reverse the decision to hire the majority candidate) and fines (the principal imposes a financial penalty for the organization whenever it discovers that the majority does not hire the most talented candidate).

(b) *Macromanagement* refers to policies based on the stock variables rather than on literature on moral wiggle room initiated by Dana et al, 2007).

Micromanagement: Information on current candidates	Macromanagement: Information on aggregate observables	Internal governance: Information on organizational practices
<ul style="list-style-type: none"> - Discretionary overrulings - Fines for discriminatory hirings 	<ul style="list-style-type: none"> - Competitive quality rewards - Affirmative action (minority quota) 	<ul style="list-style-type: none"> - Equal treatment among members - Blind and semi-blind hiring - Accountability & transparency in hiring (fines, reputations)

Figure 4: Policy interventions and principal's information.

a specific appointment: either homophily M_t (for minority quotas), or quality S_t (for competitive financial rewards based on quality).

(c) *Internal governance* refers to the set of general rules regarding internal organizational choices, for example the equal treatment of members (e.g., with respect to wages or promotions, access to resources, etc., depending on the principal's information). A second internal governance choice is the design of the hiring process. Indeed, it is widely perceived that the way candidates are selected matters: the intervention thus concerns not actual hires, but the way they are selected. The idea is either to alter the information structure somewhere along the recruitment process (blind or semi-blind hiring processes), or to harness financial or social rewards by making discrimination more salient when it happens (reputational rewards):

- *Blind and semi-blind hiring.* A principal can impose that hiring decisions be made in the absence of information about a horizontal trait (gender, race, background...). This is the case for blind auditions by orchestra (blind hiring, e.g, Goldin-Rouse 2000) or for the use of blind CVs for the pre-selection of candidates for a job (semi-blind hiring as the process' next steps are no longer blind, e.g. Bertrand-Mullainathan 2004).
- *Making discrimination in hiring more salient.* A principal (or a pressure group) can make discrimination against a talented (minority) candidate more salient and thereby mobilize social-pressure mechanisms, without exerting direct control over the decision itself. Structured interviews, anti-nepotism policies, and limits on referrals (which are known to boost homophilic recruitment by allowing the majority to discriminate by omission rather than by commission) belong to this category. Such rules can be viewed as negatively impacting the social or self-esteem of majority members who discriminate against a talented minority candidate: Structured interviews following a given script and anti-nepotism policies make discrimination more noticeable, while removing referrals limits moral wiggle room.²⁵

²⁵In contrast to blind and semi-blind hiring processes, hiring processes making discrimination more

For expositional clarity, we henceforth structure our policy analysis based on the interventions' operating channels and outcomes. We thus begin by studying interventions that we show can reduce the value of control rights (equal-treatment rules for members and minority quotas, Section 3.2), then command-and-control interventions (hiring processes and discretionary overrulings, Section 3.3), and lastly, incentive schemes (quality rewards, and fines and reputational penalties for discrimination, Section 3.4).

Objective function. The formulation of the ergodic welfare flow, $W = qS + B$, implicitly assumes that either the budget is fixed or there is no shadow cost of public funds. This assumption is natural in the case of a quota, which by definition involves no external funds. In the case of a monetary incentives for quality, one can think of a ministry of research with a fixed budget: whatever is given in terms of rewards for quality for the best institutions comes to the detriment of other institutions, who are granted a lower budget.²⁶

As announced in the introduction, the recurring themes in the following analysis are:

1. *External interventions can reduce the value of decision rights.* Majorities cannot optimize as efficiently when they face external constraints, such as equal-treatment rules for members, minority quotas, or blind hiring. This reduced value from control diminishes the appeal of entrenchment (loss-of-control-value effect), but the external constraints tend to generate inefficiencies.
2. *Interventions that make control uncertain, while still allowing the majority to fight back, such as semi-blind hiring processes and micro-management, can backfire.* The fear of an involuntary loss of control due to discretionary external interventions or restrictions on the choice of candidates encourages the current majority to build a buffer against such events, i.e., to super- or fully-entrench itself (precautionary-buffer effect).
3. *Incentives work better than command-and-control.* As in other areas of policy making, a direct intervention on the externality (here, the insufficient demand for minority talent by the majority), such as e.g., competitive quality rewards or symbolic

salient do not change the information available to the majority during the hiring process, but inflict (symbolic) penalties if the majority uses its information against minority candidates.

²⁶For the sake of exhaustivity, we consider in Online Appendix K.1 a strictly positive cost of public funds in our analysis of quality incentives, to allow for an increased total budget.

penalties for discrimination, is preferable to indirect approaches, that do exhibit side effects.

3.2 Reducing the value of control rights: Equal treatment of members and affirmative action

3.2.1 Equal treatment of members

Equal-treatment rules, as their name indicates, state that decisions should not depend on a candidate's horizontal type. In the case of equal treatment of members, such rules enforce the absence of *post-hiring* discrimination based on horizontal types: allocative discrimination restricting access to a resource for minority members, or sorting them in different tasks or career tracks (biasing assessments or promotions in favor of a given group), within-job wage disparities for equal qualification and performance, etc. We show that equal-treatment rules (weakly) foster meritocracy and thus improve quality. Notwithstanding, equal-treatment rules may entail efficiency losses – e.g., in Online Appendix G, such rules may require "locating" (on a Hotelling line) a public good at an equal distance from the two groups' respective bliss points, while the efficient location would be closer to the majority's own bliss point (especially when the majority is large and "distance costs" convex).

Proposition 5 (Equal treatment of members). *Policies that enforce equal treatment of members promote meritocracy and (weakly) increase ergodic quality (S), although not necessarily welfare (W).*

3.2.2 Affirmative action (minority quotas)

Suppose that the principal observes M_t for each t , and mandates diversity by setting a minority quota, i.e., imposing that the minority count at least R members at the end of any given period. Since it is suboptimal for the principal to impose parity,²⁷ we focus on weaker forms of affirmative action (AA) with minority quotas $R \leq k - 1$.

By limiting the majority size, the quota reduces homophily benefits. Moreover, it discriminates against the most talented candidate if at the moment of the vote, the quota binds (i.e. $M = N - R$) and the majority candidate is more talented. However, there is an indirect effect: control is less appealing both because the majority is constrained and

²⁷Suppose that the principal imposes parity (so at the end of any period the two groups are equally represented). Then, the average quality of the coopted member (\bar{x}_S) is smaller than in both the basic-entrenchment and meritocratic equilibria and homophily benefits are minimized.

because the minority is favored. That effect may lead to a “constrained meritocratic” equilibrium in which the recruitment choice is meritocratic except perhaps when $M = N - R$ at the moment of the vote, and ties are broken either in favor of the majority candidate (standard favoritism) or the minority candidate (reverse favoritism).²⁸

Proposition 6. (Affirmative action: Minority quota)

- (i) **Existence region.** AA in the form of a minority quota $R \leq k - 1$ can induce the organization to switch from basic entrenchment to constrained meritocracy.
- (ii) **Welfare.** Suppose $s/b < \rho^m$, and so, basic entrenchment prevails under laissez-faire. When the quota R is too low to induce the organization to engage in constrained meritocracy, AA reduces both ergodic quality (S) and welfare (W). When the quota induces the organization to engage in constrained meritocracy with standard favoritism, there exists a cut-off for the probability x that the minority candidate is more talented than the majority one such that AA yields a higher ergodic quality (S) than laissez-faire if and only if x is above the cutoff – the higher the quota, the lower the cut-off. AA thus increases welfare (W) if the quota is sufficiently large, x is above the cutoff and the welfare-weight (q) on quality is sufficiently large.

By reducing the value for majority members of controlling recruitments, AA can thus induce the organization to switch from basic entrenchment to meritocracy, thereby having an impact on recruitments above and beyond the immediate hiring constraint when the threshold is reached.

One can think of possible improvements to this rigid quota. For example, a “quota on average over some period” (either on the average majority or group sizes, or on the average shares of majority and minority candidates recruited) would offer the majority more flexibility to adapt to a time-varying supply of talent. However, the principal’s information may be limited as to how job-specific talent supply is affected by, say, gender-and-talent-contingent occupational choices. This may make it difficult to calibrate “fair quotas”.

²⁸When s/b is very high, the efficiency loss at $M = N - R$ becomes extremely costly and majority members may be willing to pick the minority candidate at lower majority sizes whenever the latter is as talented as the majority one in order to avoid reaching a majority size of $M = N - R$ at a later period. *Constrained meritocracy with reverse favoritism* may thus arise in equilibrium: at large majority sizes, majority members vote for their candidate if and only if he is strictly more talented than the minority candidate. How relevant is such reverse favoritism? For s/b high, meritocracy is likely to prevail in the organization and regulators unlikely to intervene on an *ad hoc* basis. But an economy-wide affirmative action rule would apply even to organizations that would otherwise be meritocratic, giving rise to reverse favoritism.

3.3 Command-and-control interventions

3.3.1 Blind and semi-blind hiring

Suppose the organization faces a pool of four candidates, two from each group. To keep things comparable, suppose that the distribution of talent among the four candidates is such that the joint distribution of the talents of each group’s most talented candidate is as described in Section 2 (portrayed by parameters x and α). And so, whether there are two or four candidates is irrelevant as long as the majority knows the candidates’ attributes (“transparent selection”). This is the case in our laissez-faire benchmark (baseline model), in which members (or HRM personnel, perfectly aligned with the majority) observe each candidate ι ’s horizontal trait, $X_\iota \in \{A, B\}$, and talent, $s_\iota \in \{0, s\}$.

The polar case of transparency is a blind process, in which the selection proceeds in the absence of any knowledge of the candidates’ horizontal attributes. Between these two processes lies the sequential, “semi-blind process”: First, a short list of two is selected in the absence of horizontal information; second, the two candidates in the short list are interviewed, and their horizontal types are revealed and the final selection is made from the short list.

We assume that while the horizontal trait (gender, race...) is costlessly observable, not learning it requires the absence of an interview that would further reveal more information about “quality”. Even in the case of a blind audition for an orchestra position, meeting with the candidate may reveal information about their future integration, their vision of the job, their desire to stay etc (in the same way a flyout for an academic position reveals information not contained in the JM package). We thus posit that a blind audition or CV reveals only a signal z_i , such that the signal z_i and the candidate’s actual talent s_i satisfy MLRP. Lastly, we assume that the probability that in the first round, the two minority candidates have strictly higher signals z than the two majority candidates is strictly positive (which is a very weak assumption).

- *Blind hiring.* The selection is made based on $\{z_\iota\}_{\iota \in \{1, \dots, 4\}}$ only.²⁹
- *Semi-blind hiring.* The first round is blind, so the pre-selection for the short list is made based on $\{z_\iota\}_{\iota \in \{1, \dots, 4\}}$ only. The final selection builds on the true characteristics $\{X_\iota, s_\iota\}$ of the remaining two candidates.

²⁹Blind hiring can be interpreted both as an intervention annihilating the value of control rights (which minority quotas and rules enforcing equal treatment of members only reduce), and as a command-and-control intervention.

The symmetry of the problem makes it optimal for the majority to select (blind) or pre-select (semi-blind) the candidate(s) with the best signal. The comparison between laissez-faire and blind hiring is straightforward: blind hiring promotes minority hiring, but may induce mistakes in estimating the vertical dimension. Blind hiring thus improves ergodic quality (S) if and only if (a) the majority is (basically) entrenched under laissez-faire, and (b) the signal is sufficiently precise.³⁰

By contrast, semi-blind hiring can induce a strategic reaction by the majority.³¹ Adapting our previous terminology, we refer to the set of equilibria that are either meritocratic, basically-entrenched or super-entrenched at level $l \geq 1$, as the set of *canonical equilibria with semi-blind hiring*.³²

Proposition 7 (Blind and semi-blind hiring). *Suppose $x < 1/2$.*

(i) **Blind hiring** induces a strictly higher ergodic quality (S) if and only if (a) the majority is (basically) entrenched under laissez-faire, and (b) the talent signal z is sufficiently precise.

(ii) **Semi-blind hiring.** Full entrenchment is the unique canonical equilibrium with semi-blind hiring for any $s/b \in [1, \rho]$, for some $\rho \in (1, \rho^m)$. Hence, for $s/b \in [1, \rho]$, semi-blind hiring reduces ergodic quality (S).

The result for semi-blind hiring illustrates a running theme of our analysis: When control is uncertain, i.e., in environments in which the majority can lose control with a strictly positive probability regardless of its strategy,³³ super-entrenchment and full-entrenchment can arise. Uncertainty about control can stem from well-meaning interventions such as a semi-blind hiring process (as in this Section), or discretionary interventions (Section 3.3.2), or alternatively from uncertain voting participation or uncertain group allegiance

³⁰Blind hiring always induces a lower ergodic aggregate homophily payoff (B) than basic entrenchment and meritocracy under laissez-faire. Hence, even when blind hiring induces a strictly higher ergodic quality (S), it induces a lower ergodic welfare (W) when the quality-over-homophily ratio, $s/b \geq 1$, and the quality externalities, $q \geq 1$, are low (close to 1).

³¹Modelling an organization searching for candidates (in the absence of homophily concerns), Fershtman-Pavan (2021) shows that if the evaluation of minority candidates is noisier than the one of majority candidates, then "soft affirmative action policies" tilting the search in favor of minority candidates can backfire and actually reduce the likelihood of a minority candidate being recruited. By contrast, in our model, semi-blind hiring only makes the "search technology" neutral, and the backfiring stems exclusively from control concerns, and not from informational frictions.

³²Specifically, in the case of semi-blind hiring, we define the meritocratic, the basic-entrenchment and the level- l entrenchment strategies as in Section 2.1.1 whenever the majority faces exactly one candidate from each horizontal group, and extend these strategies to the events in which the majority faces either two majority candidates, or two minority candidates, by assuming that in any such event, the majority picks a most talented candidate among the two. (In any such event, choosing a most talented candidate is a dominant strategy, strictly so when one candidate is strictly more talented than the other.)

³³In our baseline model, the majority can secure control (if it wants to) by being entrenched at size $M = k$.

(which we study in Section 4.4.2). In particular, all such interventions backfire for s/b close to 1 by inducing full entrenchment.

3.3.2 Discretionary overrulings

Returning to a transparent process, suppose that in each period, the majority selects among the two candidates, yet the principal can then overrule the majority and pick the losing candidate. None of the players (principal, majority, and minority) can commit. In each period, the principal learns the quality of the current candidates (or at least their quality differential) with probability λ , and remains uninformed with probability $(1 - \lambda)$. Because a non-meritocratic choice is always to the detriment of a minority candidate, it is irrelevant whether the principal observes the horizontal types of candidates.

We look for equilibria where (a) the majority is either meritocratic, basically-entrenched or super-entrenched at level $l \geq 1$, and (b) the principal overrules the majority if and only if it is informed that the majority is violating meritocracy. We refer to any such equilibrium as a *canonical equilibrium with discretionary overrulings*.

In the absence of commitment, it is an equilibrium for the principal not to intervene when it is uninformed.³⁴ Hence, for $\lambda = 0$ (always uninformed), the meritocratic and basic-entrenchment equilibria exist for the same parameter values as in the absence of intervention. At the other extreme, when $\lambda = 1$, the principal can (and will) select the best candidate in each period, and there is no real “cooptation”. Hence, let us assume that $0 < \lambda < 1$. Regardless of λ , the existence condition of a meritocratic equilibrium is unchanged, as the principal has no reason to intervene in such an equilibrium. This property however does not hold for the basic-entrenchment equilibrium. Intuitively, the possibility of intervention has two opposite effects on the principal’s welfare. When occasionally overruling the majority (which it does in equilibrium), it imposes the meritocratic choice. But the majority may become wary of losing control when $M = k$ and may thus decide to be super-entrenched so as to lower the probability of its losing control (without annihilating it completely, which is impossible). The next Proposition establishes that the ability to overrule the majority systematically backfires by generating full entrenchment for s/b close to 1.

³⁴See Online Appendix J. Intuitively, (a) from the perspective of the principal (with $q \geq 1$), the majority takes the socially optimal decision for any majority size $M \geq k + 1$, and if it is meritocratic, also when $M = k$, whereas if it is basically entrenched and tight, it takes the optimal decision with probability $1 - x \geq 1/2$; (b) if the majority is basically entrenched and tight, then its choice of candidate reveals no information on the latter’s quality to the principal, and thus a talent-blind principal cannot outperform the majority’s choice.

Proposition 8. (*Unintended effects of discretionary quality-based interventions*) Suppose $x < 1/2$.

- (i) Fix $\lambda \in (0, 1)$. Full entrenchment is the unique canonical equilibrium with discretionary overrulings for any $s/b \in [1, \rho]$, for some $\rho \in (1, \rho^m)$.
- (ii) Fix $s/b > 1$. For λ in an intermediate range (non-empty if s/b is close to 1), the principal would achieve a higher ergodic quality (S) if it could commit not to intervene.

Our result connects with the literature on dynamic conflicts with endogenous control rights – see notably Powell (2004) and Austen-Smith et al. (2019). As emphasized in that literature, the lack of commitment coupled with (sufficiently large) random shocks to the agents’ decision power lead to inefficient equilibria. In our environment, absent any intervention, the uncertainty on the identity of the departing member at the end of each period, and on the quality of future candidates contributes to (basic) entrenchment. But as Proposition 8 shows, the discretionary overrulings increase the uncertainty around majority control at all majority sizes, so much so that they can trigger full entrenchment and a less efficient equilibrium outcome.

3.4 Incentives

Finally, we consider balanced-budget (zero-sum) incentives. The latter can be provided in three manners, leading to the same outcomes.

(i) *Material rewards for talent in the organization.* Suppose that the principal (or the market) implements a quality assessment exercise after each period’s election with probability corresponding to a Poisson process of rate η . A quality assessment exercise in period t results in an end-of-period bonus accruing to the organization and shared equally among the N members. We assume without loss of generality that the bonus earned at date t is immediately paid to the organization.³⁵ For the sake of simplicity, we also assume that the bonus is linear in the number of talented members in the organization: for each talented member in the organization at the end of period t , each member receives y . Consequently, the expected incremental lifetime contribution of a new talented (relative

³⁵ Alternatively, we could have assumed that the bonus is split across several periods. Yet, frontloading the bonus is more efficient. Indeed, because members may quit, and thus $\delta \leq (N-1)/N < 1$, frontloading the bonus maximizes the incentive for good recruitment.

to mediocre) addition to each current member of the organization now writes as

$$s^+(\eta, y) \equiv s + \eta \frac{y}{1 - \delta_0(1 - 2/N)} = s \left(1 + \eta \frac{y}{\tilde{s}} \right) > s$$

while the expected lifetime utility for an incumbent member generated by the homophily payoff per new in-group member is still given by b .³⁶

Payments to organizations are financed through a general reduction of baseline transfers to (or an increase in taxes paid by) organizations. For simplicity, let us assume a continuum of the latter so that an organization is infinitesimal with regard to the overall budget, and let us look at the ergodic state. If meritocracy is an equilibrium in the absence of a quality-based bonus-malus, then it is also an equilibrium under a bonus-malus.³⁷

(ii) *Fines for unequal treatment of candidates.* Like in the discretionary overruling case, the principal occasionally observes the relative quality of the candidates. But rather than picking the losing, more deserving candidate when there is discrimination, which we show incentivizes super-entrenchment in anticipation, the principal imposes a fine on all majority members who took part in the discrimination. The principal then redistributes the proceeds either to the organization as a whole, or to minority members alone, or uses them for other goals.

(iii) *Reputational (symbolic) incentives.* When observing an unequal treatment of candidates, the principal can alternatively use the majority members' concern for reputation. Mobilizing the social-pressure and self-esteem mechanisms by making the discrimination more salient creates reputational, symbolic incentives. Structured interviews, anti-nepotism policies, and limits on referrals (which are known to boost homophilic recruitment) belong to this category. If we take the stylized view that reputations are a zero-sum game (the prestige of one comes at the expense of a disregard for others, as in models of prosocial behavior), increasing the salience of discriminatory acts implies a reputation loss for the majority when it discriminates, and an equivalent gain for the

³⁶Computations go through as in the main model with a quality-payoff-over-homophily-benefit ratio now given by s^+/b instead of s/b . Hence, for η, y sufficiently high, the ratio s^+/b is sufficiently high for the organization to reach the region where the meritocratic equilibrium exists.

³⁷Indeed, meritocracy is an equilibrium in the absence of a bonus-malus if and only if

$$s - b \geq \delta \frac{k-1}{N-1} (V_{k+1}^m - V_{k-1}^m),$$

and while a quality-based bonus-malus raises the LHS, it does not affect the RHS as the most talented candidates are always recruited under meritocracy. We refer to Online Appendix K.1 for the case of non-budget-balanced incentives.

minority (or for meritocratic majorities in other organizations).

Proposition 9 (Incentives). *Suppose $s/b < \rho^m$ (basic entrenchment prevails under laissez-faire). Any of the three zero-sum incentive schemes – material rewards for talent in the organization, fines for unequal treatment of candidates, reputational/symbolic incentives – raises ergodic quality (S) and welfare (W), strictly so if the incentive is large enough.*

3.5 Empirical evidence on policy interventions

Homophily-blind recruitments have been found to reduce discrimination: notably, identity-blind auditions improve women's hiring by orchestras (Goldin-Rouse, 2000). By contrast, and as predicted by our analysis, *semi*-blind hiring can backfire, leading to more discrimination (Behaghel et al, 2015).

Moreira-Pérez (2024) study the consequences of the 1883 Pendleton Act, which mandated exams for some employees in the largest US customs-collection districts, and find that although the act improved targeted employees' professional background, it incentivized discriminatory hiring in exam-exempted positions, which can be interpreted in the light of our model as the majority building a "safety buffer" in the administration at large. Mirroring our model's spillover logic, mandated exams selected the best where enforced, but organizations reallocated hires into exam-exempt slots to preserve control, a classic form of buffering.

Among interventions that reduce the value of control rights, our model predicts that a minority quota can induce an organization to switch from (basic) entrenchment to constrained meritocracy, and consequently lead to a higher average quality of majority members, which is consistent with the evidence in Besley et al (2017).

In line with our analysis of reputational/symbolic rewards, more structured hiring reduces in-group favoritism while improving performance (Hoffman et al, 2017). Moreover, reputational interventions, such as raising awareness of the inappropriateness of discrimination, can backfire and result in more discrimination when such interventions are perceived as short-lived, whereas policies that put in place durable practices and structures do reduce discrimination lastingly (Hirsh-Cha, 2017). Similarly, the mode of resolution of employment discrimination lawsuits influences the organization's future behavior: lawsuits that attract sufficient media coverage commit the organization to less discrimination and lead to lasting higher diversity; by contrast, lawsuits that do not attract sufficient public scrutiny do not reduce, and can even increase the discrimina-

tion against minorities within the organization (Hirsh-Cha, 2018). Consistent with our prediction, salient non-discretionary changes (court-mandated, visible to markets and media) durably improve diversity, whereas one-off payouts without credible monitoring can backfire.

4 The collegial bias against meritocracy: Further drivers

This section considers robustness results and extensions, causing organizations and their members to depart from the (meritocratic and basic-entrenchment) canons we described in Section 2. More pervasive and/or more intense forms of entrenchment arise, indicating that in practice, one can expect (significantly) fewer meritocratic recruitments than described by these canons.

4.1 The "weak link" principle in asymmetric environments

We begin with an observation on asymmetric environments – with a group having stronger homophily preferences, or a higher patience – showing that "it takes two to build meritocracy": a meritocratic equilibrium exists only if the group with the strongest homophily preferences, or the highest patience finds it optimal to behave meritocratically.

Asymmetric homophily benefits. Suppose that type-*A* agents have stronger homophily preferences than type-*B* agents, i.e. that $b_A > b_B$. So, from the point of view of meritocracy, group *A* is the "weak link".

Proposition 2'. (Asymmetric homophily benefits) *With asymmetric homophily benefits b_A, b_B such that $b_B < b_A < s$,*

- (i) *The meritocratic equilibrium exists if and only if $s/b_A \geq \rho^m$.*
- (ii) *The basic-entrenchment equilibrium exists if and only if $s/b_B \leq \rho^e$,*
- (iii) *If homophily benefits are sufficiently dissimilar (so that $\rho^e b_B < \rho^m b_A$), the unique MPE in pure strategies when $\rho^e b_B < s < \rho^m b_A$ is such that, regardless of initial conditions, type-*A* members eventually form a basically-entrenched majority.*

(iv) If $\rho^m b_A < \rho^e b_B$, the meritocratic and basic-entrenchment equilibria coexist over a non-empty range of qualities s . Over this range, the meritocratic equilibrium Pareto-dominates the basic-entrenchment one.

As Proposition 2' shows, it takes two to build meritocracy, and it takes only one to destroy it. Indeed, when facing a rival group with strong homophily preferences (A , with $s/b_A < \rho^m$), an otherwise meritocratic group (B , with $s/b_B > \rho^m$) anticipates the entrenched behavior of its rival, and thus either opts for "preemptive entrenchment" (when $s/b_B < \rho^e$), or for meritocratic recruitments (when $s/b_B > \rho^e$), only to eventually lose control to the other group, which then entrenches its majority.³⁸

Asymmetric patience. The same *weak-link principle* holds when the two horizontal groups have different pure-time preferences $\delta_{0,A}, \delta_{0,B}$ (and similar homophily preferences). It is then the most patient group who triggers entrenchment.

Suppose for instance that $\delta_{0,A} > \delta_{0,B}$. Hence, the meritocratic equilibrium exists if and only if $s/b \geq \rho^m(\delta_{0,A})$, while the basic-entrenchment equilibrium exists if and only if $s/b \leq \rho^e(\delta_{0,B})$.³⁹ If $\rho^e(\delta_{0,B}) < \rho^m(\delta_{0,A})$, the unique MPE in pure strategies when $\rho^e(\delta_{0,B}) < s/b < \rho^m(\delta_{0,A})$ is such that type- A members eventually form a (basically) entrenched majority.

4.2 A continuum of vertical types

We have assumed so far that talent can take only two values. When talent is smoothly distributed in \mathbb{R}_+ , for the natural generalization of canonical equilibria developed below, full meritocracy never prevails, as the majority always prefers an in-group candidate over a slightly more talented out-group candidate. But, as we will see, we can still order equilibria in terms of their "level of meritocracy". Our previous insights generalize: (i) a stronger majority engages in more meritocratic recruitments, and (ii) whenever several equilibria coexist, they can be ranked from more to less meritocratic and Pareto-compared.

Generalizing canonical equilibria to arbitrary talent distributions, equilibria can be

³⁸Put differently, an increase in b_A can have long-term consequences on the organization's dynamics, whereas an increase in b_B (still below b_A) has at most short-term consequences.

³⁹We refer to $\rho^m(\delta_0)$ and $\rho^e(\delta_0)$ as the cutoffs that obtain from the no-profitable-deviation condition for a group with pure-time preference δ_0 (see expressions in Online Appendix C.2.1). [Note that the cutoffs $\rho^m(\delta_0)$ and $\rho^e(\delta_0)$ for a given group do not depend on the other group's time preference: they depend only on the group's own time preference (δ_0), the talent distribution (parameters x, α) and the transition probabilities (driven by the two groups' strategies).] By Proposition 2.(iv), we know that $\rho^m(\delta_0)$ and $\rho^e(\delta_0)$ increase with δ_0 .

described as a sequence of strictly positive cut-offs $(\Delta_M)_{M \in \{k, \dots, N-1\}}$ such that a majority of size M recruits the out-group candidate with (discounted) talent \hat{s} against the in-group candidate with (discounted) talent s if and only if $\hat{s} - s > \Delta_M$. We show in Online Appendix M that in any such equilibrium, $\Delta_M > b$ for any $M \in \{k, \dots, N-1\}$. Intuitively, in-group recruiting when $b > \hat{s} - s$ yields a double dividend – a larger homophily payoff and a tighter grip on the organization –, and thus for a minority candidate to be considered by the majority, her talent must exceed the majority candidate's by strictly more than the homophily benefit: $\hat{s} - s > b$.

We denote by \prec the order relation defined over the set of decision rules such that $\Delta \prec \Delta'$ if and only if $\Delta_M < \Delta'_M$ for all $M \in \{k, \dots, N-1\}$. We will then say that the former decision rule is more meritocratic.

Definition. Let \mathcal{G} be the set of continuous joint distributions of (s, \hat{s}) , i.e. resp. the quality of the majority and the minority candidate, with support in $[0, +\infty)^2$ such that $\mathbb{E}[\max(\hat{s}, s + b)] < \infty$, and $(\hat{s} - s)$ is symmetrically distributed around 0 with $\mathbb{P}(\hat{s} - s > b) > 0$ and such that, letting the function h be defined by

$$h(\Delta) \equiv \mathbb{E}[(s + \Delta)\mathbf{1}\{\hat{s} - s \leq \Delta\}] + \mathbb{E}[\hat{s}\mathbf{1}\{\hat{s} - s > \Delta\}],$$

the functions $\Delta \mapsto [h(\Delta) - \Delta/2]$ and $\Delta \mapsto [\Delta - h(\Delta)]$ are strictly increasing with $\Delta \in (b, \bar{s})$ where $\bar{s} = \sup(\hat{s} - s) \in (b, +\infty]$.

In equilibrium, the talent bar Δ faced by minority candidates equals the homophily and control-driven benefits that a majority candidate brings to majority members. The function h thus captures the sum of the expected benefit from the current hire (\hat{s} when the minority candidate is recruited, $s + \Delta$ when the majority one is).⁴⁰

The set \mathcal{G} notably includes the set of (full support) continuous joint symmetric distributions with finite-mean marginals. It also includes the case where the majority candidate has a fixed talent $s \geq 0$ and the minority candidate a talent $s + D$ where D is a (full support) random variable with a continuously differentiable distribution over $(-s, s)$ symmetric around 0.

Proposition 10. (A continuum of vertical types) Assume talents are distributed

⁴⁰The restrictions on h for distributions in \mathcal{G} have an intuitive interpretation: raising the homophily and control-driven benefits that a majority candidate brings to the majority members by a small amount $d\Delta$ must induce a gain that is (i) strictly higher than $d\Delta/2$ if majority candidates are recruited more often than minority ones (and such gains are thus collected with a probability higher than 1/2), but (ii) strictly lower than $d\Delta$ if minority candidates are still recruited with a strictly positive probability (and such gains are thus collected with a probability lower than 1).

according to a joint distribution $G \in \mathcal{G}$. Any symmetric MPE described by a sequence of cut-offs $(\Delta_M)_{M \in \{k, \dots, N-1\}}$ is such that $\Delta_M > b$ for any M , and the sequence $(\Delta_M)_M$ is strictly decreasing: a stronger majority discriminates less than a weaker majority.

Moreover, whenever they coexist, any two such equilibria with distinct decision rules Δ and Δ' , can be ranked by the order relation \prec . If $\Delta \prec \Delta'$, then the equilibrium characterized by the decision rule Δ (which is more meritocratic than the one described by Δ') is preferred at any majority size by all current majority members, and for δ_0 small, by all current minority members as well.

Consistently with our previous stance, let us consider the following equilibrium selection: At any time t , the current majority can choose its preferred continuation equilibrium. In the binary-talent case, Proposition 3 implies that this criterion selects the meritocratic equilibrium whenever it exists – as we assume throughout our analysis of the binary case. Correspondingly, in the continuous-talent case, Proposition 10 thus implies that, whenever several equilibria coexist, the most meritocratic one (i.e., with the lowest sequence of cut-offs Δ) is selected.⁴¹

4.3 Generalization to non-linear homophily benefits

Our analysis can be extended to non-linear homophily benefits. Convex homophily benefits arise for instance when facilities or regulations must be added to accommodate the existence of a minority, or when a group's reaching a critical size delivers additional opportunities to its members, e.g., because of supermajority clauses for some decisions. Conversely, concave homophily benefits arise if there are decreasing returns to having one more in-group member (e.g., limited time for "horizontal" interactions) or increasing returns to having one more out-group member (e.g., benefits from diversity).

We show in Online Appendix N that concave homophily benefits still give rise to equilibria that are either meritocratic, basically-entrenched or super-entrenched, whereas convex homophily benefits can induce equilibria in which all recruitments are entrenched above a certain majority size.⁴² In other words, while with linear (as in our baseline

⁴¹For any distribution in \mathcal{G} , a most meritocratic equilibrium exists. Indeed, Proposition 10 implies that any two symmetric MPE described by distinct cut-off rules Δ and Δ' , can be ranked by the order relation \prec . Moreover, all cut-off rules are bounded from below by the constant rule with $\Delta_M = b$ for all M . Suppose by contradiction that the set of equilibrium cut-off rules has no minimum (according to \prec). The above arguments imply that this set has a finite infimum. By continuity, the equilibrium conditions still hold at the infimum, and so the set of equilibrium cut-off rules has a minimum, i.e., a most meritocratic equilibrium.

⁴²Starting from a majority size above the threshold, such equilibria thus generate the same (on-path) dynamics as full entrenchment.

model) or concave homophily benefits, our model predicts that larger majorities discriminate less than thinner ones, the opposite may hold with (sufficiently) convex homophily benefits: then, larger majorities may actually discriminate more than thinner ones.

4.4 Further drivers of super-entrenchment under laissez-faire

Returning to the binary-talent case, the most obvious case yielding super-entrenchment is $s \leq b$, which trivially leads to full entrenchment. Section 4.3 noted that, even for $s > b$, non-linear homophily benefits may lead to super-entrenchment. Let us now describe two other drivers of super- and full-entrenchment under laissez-faire: homogamic evaluation capability and uncertain voting participation or identification of group allegiance.

4.4.1 Homogamic evaluation capability

We assumed that all members are equally proficient at evaluating the talents of in- and out-group candidates. However, some environments exhibit an asymmetry in this ability. For example, econometricians are better placed than development economists to evaluate an econometrician, and conversely.

When only in-group evaluation is feasible, the majority still selects the majority candidate if the latter has quality s . So, the minority candidate is only considered when the majority candidate has quality 0. With $\bar{x} \equiv x + (1 - 2x)\alpha$ denoting the probability of a given group's candidate being of quality s , the conditional quality of the minority candidate is then

$$s^\dagger \equiv \frac{x}{1 - \bar{x}}s \leq s.$$

Let us focus on the case of "pessimistic expectations" (or negative stereotypes): $s^\dagger \leq b$. This case arises when correlation is high (x low) and average quality low (\bar{x} low), so the majority is pessimistic about the minority candidate's talent when its own candidate lacks talent. [Departing from the Bayesian framework, this case would also be more likely if the majority members had a negative stereotype about minority members' talent.] When $s^\dagger \leq b$, talented minority candidates are victims of asymmetric information as their talent and possibly superiority over untalented majority candidates is not observed by the majority. Therefore, the majority is fully entrenched: it keeps admitting solely majority candidates and ends up being homogeneous. This implies that imperfect information (in the form of homogamic evaluation capability) may transform a basically-entrenched or

meritocratic organization into a fully entrenched one.⁴³

Proposition 11. (*Homogamic evaluation capability*) *If $s^\dagger \leq b$ (pessimistic expectations/negative stereotype), the majority coopts only candidates of the in-group and thus becomes homogeneous. Homogamic evaluation capability then lowers the ergodic welfare relative to perfect information.*

Remark: Disclosing vertical information. Proposition 11 suggests a policy intervention when $s^\dagger \leq b$ and the principal has hard information about candidates' talent (e.g., via test scores): disclosing such information improves ergodic welfare – consistently with the evidence in Hoffman et al. (2017), who show that under discretionary hiring, the availability of test scores raises the quality of appointments (as measured by subsequent job tenure).⁴⁴

4.4.2 Uncertain voting participation or identification of group allegiance

We have assumed so far that all members of the organization vote, and that they vote as expected. Relaxing one or the other assumption, we note that absenteeism (whether due to illness or alternative obligations) or uncertain identification of group allegiance (if candidates are able to masquerade as belonging to the majority group, or if homophily-vs-quality preferences are heterogeneous within groups and privately observed, with some individuals putting a much higher weight on talent relative to homophily) may incentivize groups to secure majorities of more than one vote so as to minimize the probability of a control switch. Consequently, with uncertain voting participation or identification of group allegiance, even large majorities may find it optimal to stand in the way of talented minority candidates.

We focus on uncertain voting participation, referring to Online Appendix P for the imperfect identification of group allegiance. For any majority size $M \in \{k, \dots, N-1\}$, let

⁴³Online Appendix O studies the more complex case in which the minority candidate is in expectation *preferred to* an untalented majority candidates ($s^\dagger > b$). Then, analogues of the meritocratic and basic-entrenchment equilibria exist, in which the minority candidate is given the benefit of the doubt except perhaps when the majority is tight ($M = k$). We show that such meritocratic and basic-entrenchment equilibria with homogamic evaluation capability yield a lower ergodic aggregate welfare than their perfect-information counterparts.

⁴⁴By contrast, in the absence of hard information and of commitment, communication in the form of one-shot cheap talk cannot operate due to a form of winner's curse. Because the majority picks its candidate whenever talented, the minority infers that whatever message it sends can only have an impact when the majority candidate is untalented. But conditional on a low-quality majority candidate, the minority always prefers its own candidate, and so any message sent by the minority is necessarily uninformative.

$\Lambda(M)$ be the probability that, because of absenteeism, the minority's choice prevails.⁴⁵ We assume that the majority is strictly more likely than the minority to win the vote, and the more so, the greater the majority size, with the majority being certain to win for sufficiently large majority sizes (clearly so for $N - 1$):

$$\begin{cases} \Lambda \text{ decreases with respect to majority size } M, \\ \Lambda(M) \in (0, 1/2) \text{ for any } M \in \{k, \dots, k + l - 1\}, \quad \text{and} \quad \Lambda(M) = 0 \text{ for any } M \geq k + l \end{cases} \quad (2)$$

While the Λ function can capture correlation in absenteeism, either within groups or across the entire population of members, an interesting case occurs when absences are i.i.d. (the Bernoulli case). That case satisfies (2) with $\Lambda(M) > 0$ for all $M < N - 1$.

We look for pure-strategy symmetric MPEs in which a stronger majority makes (weakly) more meritocratic recruitments. In contrast to the baseline model, the minority's strategy now matters at any majority size M at which $\Lambda(M) > 0$. When looking for level- l super-entrenchment equilibria, we now look for equilibria in which (a) the majority is super-entrenched up to level l , and (b) the minority always votes for its in-group candidate whenever it is pivotal with a strictly positive probability, i.e., whenever $M \leq k + l - 1$.

Proposition 12. (*Uncertain voting participation and super-entrenchment*) *Let Λ satisfy (2) and $x < 1/2$. For s/b sufficiently close to 1, super-entrenchment at level l is the unique pure symmetric MPE such that a stronger majority makes (weakly) more meritocratic recruitments. In particular, if $l = k - 1$ as in the Bernoulli case, the possibility of absenteeism triggers full entrenchment for s/b sufficiently close to 1.*

When Λ satisfies (2) with $l < k - 1$, the majority is "safe" at any majority size $M \geq k + l + 1$ as it will still control the outcome with probability 1 in the next period. Therefore, meritocratic recruitments are optimal at these majority sizes.

Remark: Fighting absenteeism. Interestingly, under the conditions of Proposition 12, strong in-group discipline, inasmuch as it reduces voting uncertainty, makes hiring more meritocratic. More generally, any policy intervention curbing absenteeism may make all

⁴⁵We thus assume that absenteeism in a given period is independent of the candidates' qualities in that given period. While we allow for a wide range of absenteeism functions (in particular as we allow for correlation in voting turnout), condition (2) below may not be warranted if voting participation is strategic rather than caused by exogenous events.

members better off by reducing the degree of entrenchment.

4.5 Anterooms for appointments

We have so far viewed the appointment process as an organizational choice between recruiting a candidate and letting them go away for good. While a first step, this assumption ignores the possibility that appointments may result from a dynamic process operating outside or inside the organization. First, turned-away candidates may be persistent and later reapply. Second, the organization may groom junior members for possible promotion to senior positions. This section analyzes these two possibilities, which display several similarities.

4.5.1 Candidates can re-apply

We investigate the consequences of unselected candidates being able to re-apply. Unsuccessful candidates keep re-applying until they are recruited.⁴⁶ For the sake of exposition, we make a further simplifying assumption: $\alpha = 0$, so that in any period, the new majority and minority candidates are equally talented if and only if they both are untalented (which happens with probability $(1-2x)$), and the unconditional probability that a new candidate is talented is given by x . This assumption implies that under meritocratic hiring, talented candidates are always immediately hired and so the ability to re-apply is irrelevant on an equilibrium path. However, the knowledge that talented minority candidates will reapply lowers the cost of entrenchment and thus favors discrimination.

Proposition 13. (*Reapplying for membership*) *Assume $\alpha = 0$. Basic entrenchment yields the majority a higher value function when candidates reapply than when they cannot: being able to "keep in store" a talented minority candidate when the majority is tight reduces the cost for the majority of turning down her application. Moreover, the existence region for the meritocratic equilibrium shrinks when the organization can store applications.*

4.5.2 Hierarchies and the glass ceiling

The expression "glass ceiling" refers to the difficulty for women (or minorities) to rise beyond a certain level in a hierarchy. While there are various hypotheses for its existence, whose relevance is reviewed e.g., in Bertrand (2018), we here investigate whether the

⁴⁶Our results would still hold if we assumed instead that such candidates stopped re-applying following some Poisson process.

desire for the dominant group to retain control might be a factor.

Online Appendix R considers a (large) two-tier organization with (many) senior and junior positions. At each point in time, a fraction of seniors exogenously departs and is replaced by juniors promoted to seniority. A fraction of juniors exogenously quit the organization as well. Flows out of the junior pool are offset by new recruitments. Seniors have control over hiring and promotion decisions.

We say that a glass ceiling exists if the probability of promotion of a talented majority member is higher than the one of a minority member. Even if majority dominance and favoritism contribute to hiring discrimination against minorities, it is not a priori obvious that they imply a lower rate of promotion for the latter within the organization. Indeed, hiring discrimination implies that minority recruits are fewer and more talented than majority ones. Assuming that hiring and promotions are controlled at the senior level, we nonetheless show that a glass ceiling arises for two reasons:

- *Concern for control*: as earlier in the paper, control allows and incentivizes groups to engage in favoritism. This in turn implies some discrimination in promotions, which in general exceeds that at the hiring level (if any).
- *Differential mingling effect*: for organizational reasons, senior members tend to hang around more with senior members than with junior ones. Their homophily concerns are therefore higher for promotions than for hiring decisions.⁴⁷

Proposition 14. (*Glass ceiling*) *In the hierarchical organization's steady state, hiring at the junior level is fully meritocratic. By contrast, there exists a glass ceiling for minority juniors: A talented minority junior is less likely to be promoted than a talented majority junior.*

Our version of the glass ceiling paradigm suggests a new answer to the “macro puzzle” regarding the ubiquitous control by men of the upper echelons. While at the micro level homophily explains why a majority can be self-perpetuating in the long run, it does not address the possibility that some firms could take advantage of the resulting wasted female talent. The standard explanation for this is the statistical discrimination theory, which presumes that individual talent is imperfectly observed: The minority does not invest in

⁴⁷In line with Carmichael (1988) and Friebel-Raith (2004), it may thus be optimal for the seniors' majority not to let current juniors coopt new juniors as a majority of out-group juniors may engage in un-meritocratic hiring in order to increase their chances of being appointed to the senior board. This optimality result may not hold if for instance, juniors are better able than seniors at scouting talented candidates.

human capital in anticipation of future discrimination. A different hypothesis is provided here: The minority’s talents instead are broadly used by incumbent organizations, but few make it to the top due to the concern-for-control and differential-mingling effects.

5 Related theoretical literature

Discrimination theory. It shares with the literature on the economics of discrimination initiated by Becker (1957) the idea that homophily may lead organizations to disfavor minority members in their hiring decisions. Becker, though, famously emphasized that competitive market forces may make such discrimination vacuous, while we look at organizations facing imperfect market pressure. Also, Becker’s analysis is static while the focus of our study is on the evolution of the organization. In thinking about policies that protect minorities, our work is akin to the extensive literature on affirmative action (see Fryer-Loury 2005 for an overview). In Coate-Loury (1993), employers have a taste for discrimination and a principal wants to boost minority workers’ incentives to invest in skills. Affirmative action gives the minority prospects and, if modest, boosts its incentives, but if extensive, creates a “patronizing equilibrium” and reduces incentives. In Rosen (1997)’s statistical discrimination model, a group of workers who find it hard to get a job in competition with candidates from the outgroup become less choosy; they apply for jobs for which they are less suited, and knowing this, firms rationally discriminate against group members and in favor of the outgroup.

Recruiting like-minded candidates. Our emphasis on cooptation is reminiscent of the theories of clubs (initiated by Buchanan 1965) and of local public goods (e.g., Tiebout 1956, Jehiel-Scotchmer 1997). A couple of contributions examine the dynamics of organizational membership assuming, as we do, that current members think through the impact of joiners on future recruitment decisions. They consider contexts rather different from ours, though. In particular, they stress the time variation of the size of the organization. Barberà et al. (2001) look at clubs in which each member can bring on board any candidate without the assent of other members. They are interested in the forces that determine the growth or the stagnation of organizations. A member’s (unilateral) decision of coopting a candidate hinges on the number of additional candidates whom the newly admitted one brings in the future; for instance, a member may not vote for his friend, because his friend may bring enemies to the group. Roberts (2015), like us, assumes majority rule, but posits that individuals care only about the (endogenous) size

of the organization; there is a well-determined order of cooptation, with new members being more favorable to expansion than previous ones and therefore, if admitted, taking incumbent members into dynamics they may not wish. Acemoglu et al. (2012) also looks at the long-term consequences of reforms that benefit the rulers in the short run, but may imply a transfer of control in the future; for instance, a controlling elite may not want to liberalize (give political or religious rights to other citizens) by fear of a slippery slope that would later entail a loss of control.⁴⁸

Recruiting talent under incomplete information. Section 4.4.1 on homogamic evaluation capability bears resemblance with Board et al. (2025), which assumes that talented people are better at identifying new talents, hence deriving rich dynamics. Section 4.4.1 also considers homogamic evaluation capability, but in the horizontal dimension rather than the vertical one; there may then be a separation between information and control, unlike in Board et al.⁴⁹

Trade-off between talent and like-mindedness. Cai et al. (2018) analyzes the dynamics of a three-member club. Players are characterized by a vertical and a horizontal type, but unlike in our paper, homophily benefits are constant-sum (they stand for the sharing of spoils), while they are not in our model. Sections 2 and 4.2 generalize the analysis of Cai et al. to an arbitrary-size organization, arbitrary homophily benefits and a larger set of talent distributions, deriving new insights. While Cai et al.’s model includes costly search for candidates,⁵⁰ our model allows for a much larger scope of inquiry. Notably,

⁴⁸Acemoglu et al. (2018) emphasize that when individuals anticipate to benefit from upward social mobility in the future (i.e., to join the ranks of the elite), they may exhibit more conservative behavior in the present – and conversely when they expect to suffer from downward social mobility. Compared with our study of uncertain voting participation/group allegiance (Section 4.4.2), in Acemoglu et al (2018), individuals expect their own preferences to change in the future (as they become richer/poorer), whereas in our model, individuals expect their colleagues’ preferences to change in the future.

⁴⁹Moldovanu-Shi (2013) model also exhibits heterogeneous evaluation capabilities. Members of a committee sequentially assessing candidates for a job and coopting using the unanimity rule each have a superior expertise in evaluating a candidate’s performance along the dimension he cares most about. The focus is on the acceptance standards and the comparison between a dictator and a committee; given the focus on a single job opening, the dynamics of control are not investigated. In Egorov-Polborn (2011), similar backgrounds (homophily dimension) facilitate the estimation of others’ ability. A force pushing toward homogeneity of organizations is then the winner’s curse: competition among employers makes it more likely that organizations will hire majority candidates, on whom they have superior information.

⁵⁰An interesting insight of their analysis that is not (but could be) present in our model is the possibility of “intertemporal free riding”: Even in a homogenous population (which corresponds to $b = 0$ in our model), current members will not maximize social welfare; for, in Cai et al., members engage in costly search for candidates and as current members are not infinitely lived and thus do not enjoy the benefits of quality recruitment as long as the organization, they underinvest in search. A similar effect is present in Schmeiser (2012), who analyses the dynamics of board composition and the potential benefits of outside-directors rules and nominating committee regulations. In his paper, even outside directors may not stand for shareholders’ best interests, even if they can be ascertained to have no connection with insiders. The point is that, in the absence of delayed compensation, outside directors favor immediate benefits due to their limited tenure.

in contrast with Cai et al., we investigate super-entrenchment and explore its drivers – in particular, we show that super-entrenchment can stem from non-linear homophily benefits, uncertain voting participation or homogamic evaluation capability, or else be the unintended consequence of several policy interventions. Moreover, while Cai et al. focus on finding the optimal voting rule in a three-member club,⁵¹ we study a distinct and wide set of familiar policy interventions, including affirmative action, quality-based rewards, discretionary overrulings of majority appointments, curbing absenteeism, etc. In particular, we describe how such policies generate two conflicting effects: the *loss-of-control-value* effect and the *precautionary-buffer* effect. Finally, interesting extensions to more than two groups – see the conclusion – require studying N larger than 4.

Glass ceiling. In Athey et al. (2000), players also have a horizontal (gender) and vertical (talent) types. Ability to fill a senior position depends on intrinsic talent and on mentoring received as a junior member. Mentoring is type-based, and so majority juniors receive more mentoring and are favored in promotions. The upper level may therefore become homogenous. The organizations however may (depending on the mentoring technology’s concavity) want to bias the promotion decision in favor of minority juniors, so as to create diversity and more efficient mentoring. Control is not a focus of their paper, unlike ours.

6 Conclusion

This paper studies homophily-induced control concerns in collegial organizations. It provides rich and testable insights as to where and when such concerns lead to violations of meritocracy. It investigates several potential remedies, identifying conditions for their effectiveness and warning about their (possibly dramatic) unintended consequences.

On the positive front, this paper’s insights belong to two main themes. Firstly, meritocracy is at risk whenever control stakes are high. Distrust of the outgroup jeopardizes meritocracy as each group is more eager to cling to power if it suspects the other group would not fulfil its part of the meritocratic deal. Relatedly, an organization is only

⁵¹Our model also allows for a general investigation of voting rules in clubs of arbitrary size. Consider for instance supermajority voting rules. Suppose that a (completely uninformed) principal mandates that, to be elected, a candidate must receive at least $k + l$ votes, where $l \geq 1$. If no candidate reaches the election threshold, the principal picks one among the two candidates at random. As intuitive from our analysis, such a supermajority voting rule jeopardizes the majority’s control when it has a size below the threshold. Unsurprisingly, it can be shown that for $x < 1/2$, for s/b sufficiently close to 1, super-entrenchment at level l is the unique symmetric MPE such that a stronger majority makes (weakly) more meritocratic recruitments.

as meritocratic as its less-meritocracy-prone group (its weakest link). Similarly, longer tenures within the organization heighten control stakes and foster entrenchment. Secondly, meritocracy is at risk whenever control itself is at stake, i.e., either frail, uncertain or impeded. For instance, larger majorities are more meritocratic than thinner ones (for linear or concave homophily benefits); majorities with stronger group discipline are more meritocratic than majorities unable to prevent absenteeism or turncoats; majorities that always have the option of recruiting a majority candidate may be more meritocratic than majorities that may face no majority candidate in some periods (e.g., due to a semi-blind hiring process); independent majorities are more meritocratic than majorities exposed to outside overrulings.

On the normative front, (direct or strategic) side effects may hurt society and even the minority the policy intends to protect. Common interventions can be divided into two categories: interventions reducing the value of control rights (equal treatment of members, affirmative action, and homophily-blind selection) and those that tilt the current decision toward meritocracy (semi-blind recruitment, and micromanagement, for the command-and-control approach; penalizing hiring discrimination through material or symbolic rewards, or competitive rewards for aggregate quality, for the incentive approach).

Reducing the value of control rights introduces inefficiencies by limiting the majority's managerial capability. Tilting the current decision toward meritocracy, when of the command-and-control type, ends up going against meritocracy, as the dominant group plays cat-and-mouse with the social planner and stuffs the organization with its candidates much more than it would have done in the absence of public intervention (building a safety buffer). In the end, zero-sum interventions on incentives unambiguously foster meritocracy and improve welfare. Pigovian incentive-based policies align the majority's interests with society's, without the decision-inefficiency or buffer-building side effects inherent in the other policies.

Theoretical questions. An important extension of the research agenda is the generalization to more than two groups. The core insights from the two-group analysis – that meritocracy is threatened whenever control stakes are high and control itself is contested – continue to apply, along with their implications for meritocracy. At the same time, a richer set of new phenomena emerges. In particular, entrenchment no longer requires holding an outright majority. Instead, it can arise from the ability to form temporary or durable alliances that prevent a transfer of control to a common opponent. In this

setting, a moderate group may rely on an alliance of convenience with the smallest flank to block a takeover by the largest flank. Online Appendix T confirms these intuitions in a three-group polity (a center and two flanks) with single-peaked preferences. As in the two-group case, the dynamic Condorcet winner features full entrenchment when s/b is strictly less than 1, basic entrenchment when s/b is strictly higher than but close to 1, and meritocracy when s/b is large. Finally, as noted above, basic entrenchment requires a majority on its own for a flank, but not for the moderate group (who only needs to be part of a majoritarian alliance with a flank).

Empirical questions. We have observed that our theoretical conclusions are broadly consistent with evidence, both under laissez-faire and under policy interventions. In future investigation, the model could be tested from its basic assumptions to its predictions. For instance, the homophily incentive b has in recent years increased in some dimensions (e.g., political polarization) and decreased in others (as when social norms penalize a lack of diversity). Depending on factors such as initial conditions, the nature of internal interactions, the size of the organization⁵² or the competitiveness of the market for talent, this evolution should impact dependent variables such as the quality of recruitments and the heterogeneity within and across organizations. Does patience (e.g., longer-term perspectives for members within the organization) foster entrenchment as the model predicts? For example, the model's predictions on the role of patience may be particularly relevant when applied to local communities. People with low prospects of ever leaving a region or a neighborhood (the "somewheres", to borrow from Goodhart 2017), should be expected to be more inclined to entrench themselves, i.e., be opposed to a large immigration that would make them become a minority, while by contrast, highly mobile individuals (the "anywheres") should be more tolerant/less sensitive. We leave these exciting theoretical and empirical questions for future research.

References

- Acemoglu, D., Egorov, G., and Sonin, K. (2009). "Equilibrium Refinement in Dynamic Voting Games". mimeo.
- Acemoglu, D., Egorov, G., and Sonin, K. (2012). "Dynamics and Stability of Constitutions, Coalitions, and Clubs". *American Economic Review*, 102(4):1446–1476.

⁵²Our model indeed predicts that at least for δ_0 small, small organizations are more likely to be in the meritocratic equilibrium, yet it also predicts that conditional on entrenchment, larger organizations make on average more meritocratic recruitments.

Acemoglu, D., Egorov, G., and Sonin, K. (2017). "Social Mobility and Stability of Democracy: Reevaluating De Tocqueville". *Quarterly Journal of Economics*, 133(2):1041–1105.

Acemoglu, D. and Robinson, J. A. (2000). "Why Did the West Extend the Franchise? Democracy, Inequality, and Growth in Historical Perspective". *Quarterly Journal of Economics*, 115(4):1167–1199.

Athey, S., Avery, C., and Zemsky, P. (2000). "Mentoring and Diversity". *American Economic Review*, 90(4):765–786.

Austen-Smith, D., Dziuda, W., Harstad, B., and Looper, A. (2019). "Gridlock and Inefficient Policy Instruments". *Theoretical Economics*, 14(4):1483–1534.

Bagues, M. and Esteve-Volart, B. (2010). "Can Gender Parity Break the Glass Ceiling? Evidence from a Repeated Randomized Experiment". *Review of Economic Studies*, 77(4):1301–1328.

Bagues, M., Sylos-Labini, M., and Zinovyeva, N. (2017). "Does the Gender Composition of Scientific Committees Matter?". *American Economic Review*, 107(4):1207–38.

Barberà, S., Maschler, M., and Shalev, J. (2001). "Voting for Voters: A Model of Electoral Evolution". *Games and Economic Behavior*, 37(1):40–78.

Becker, G. (1957). *The Economics of Discrimination*. University of Chicago Press.

Behaghel, L., Crépon, B., and Le Barbanchon, T. (2015). "Unintended Effects of Anonymous Résumés". *American Economic Journal: Applied Economics*, 7(3):1–27.

Behrens, A., Uggen, C., and Manza, J. (2003). "Ballot Manipulation and the "Menace of Negro Domination": Racial Threat and Felon Disenfranchisement in the United States, 1850–2002". *American Journal of Sociology*, 109(3):559–605.

Bertocchi, G. and Spagat, M. (2001). "The Politics of Co-optation". *Journal of Comparative Economics*, 29(4):591–607.

Bertrand, M. (2018). "Coase Lecture – The Glass Ceiling". *Economica*, 85(338):205–231.

Bertrand, M. and Mullainathan, S. (2004). "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination". *American Economic Review*, 94(4):991–1013.

- Besley, T., Folke, O., Persson, T., and Rickne, J. (2017). "Gender Quotas and the Crisis of the Mediocre Man: Theory and Evidence from Sweden". *American Economic Review*, 107(8):2204–42.
- Board, S., Meyer-ter-Vehn, M., and Sadzik, T. (2025). "Recruiting Talent". *Review of Economic Studies*, forthcoming.
- Buchanan, J. M. (1965). "An Economic Theory of Clubs". *Economica*, 32(125):1–14.
- Cai, H., Feng, H., and Weng, X. (2018). "A Theory of Organizational Dynamics: Internal Politics and Efficiency". *American Economic Journal: Microeconomics*, 10(4):94–130.
- Carmichael, L. (1988). "Incentives in Academics: Why is There Tenure?". *Journal of Political Economy*, 96(3):453–472.
- Coate, S. and Loury, G. (1993). "Will Affirmative-Action Policies Eliminate Negative Stereotypes?". *American Economic Review*, 83(5):1220–1240.
- Dana, J., Weber, R. A., and Kuang, J. X. (2007). "Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness". *Economic Theory*, 33(1):67–80.
- Dustmann, C., Vasiljeva, K., and Piil Damm, A. (2019). "Refugee Migration and Electoral Outcomes". *Review of Economic Studies*, 86(5):2035–2091.
- Egorov, G. and Polborn, M. (2011). "An Informational Theory of Homophily". mimeo.
- Enos, R. D. (2014). "Causal Effect of Intergroup Contact on Exclusionary Attitudes". *Proceedings of the National Academy of Sciences*, 111(10):3699–3704.
- Enos, R. D. (2016). "What the Demolition of Public Housing Teaches Us about the Impact of Racial Threat on Political Behavior". *American Journal of Political Science*, 60(1):123–142.
- Fershtman, D. and Pavan, A. (2021). "'Soft' Affirmative Action and Minority Recruitment". *American Economic Review: Insights*, 3(1):1–18.
- Friebel, G. and Raith, M. (2004). "Abuse of Authority and Hierarchical Communication". *RAND Journal of Economics*, 35(2):224–244.
- Fryer, R. G. J. and Loury, G. C. (2005). "Affirmative Action and Its Mythology". *Journal of Economic Perspectives*, 19(3):147–162.

- Goldin, C. and Rouse, C. (2000). "Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians". *American Economic Review*, 90(4):715–741.
- Goodhart, D. (2017). *The Road to Somewhere: The Populist Revolt and the Future of Politics*. C. Hurst & Co.
- Hirsh, E. and Cha, Y. (2017). "Mandating Change: The Impact of Court-Ordered Policy Changes on Managerial Diversity". *ILR Review*, 70(1):42–72.
- Hirsh, E. and Cha, Y. (2018). "For Law and Markets: Employment Discrimination Lawsuits, Market Performance, and Managerial Diversity". *American Journal of Sociology*, 123(4):1117–1160.
- Hoffman, M., Kahn, L. B., and Li, D. (2017). "Discretion in Hiring". *Quarterly Journal of Economics*, 133(2):765–800.
- Hopkins, D. J. (2010). "Politicized Places: Explaining Where and When Immigrants Provoke Local Opposition". *American Political Science Review*, 104(1):40–60.
- Jehiel, P. and Scotchmer, S. (1997). "Free Mobility and the Optimal Number of Jurisdictions". *Annals of Economics and Statistics*, (45):219–231.
- Maskin, E. and Tirole, J. (2001). "Markov Perfect Equilibrium, I: Observable Actions". *Journal of Economic Theory*, 100:191–219.
- Moldovanu, B. and Shi, X. (2013). "Specialization and Partisanship in Committee Search". *Theoretical Economics*, 8:751–774.
- Moreira, D. and Pérez, S. (2024). "Civil Service Exams and Organizational Performance: Evidence from the Pendleton Act". *American Economic Journal: Applied Economics*, 16(3):250–291.
- Powell, R. (2004). "The Inefficient Use of Power: Costly Conflict with Complete Information". *American Political Science Review*, 98(2):231–241.
- Reny, T. T. and Newman, B. J. (2018). "Protecting the Right to Discriminate: The Second Great Migration and Racial Threat in the American West". *American Political Science Review*, 112(4):1104–1110.
- Rivera, L. A. (2012). "Hiring as Cultural Matching: The Case of Elite Professional Service Firms". *American Sociological Review*, 77(6):999–1022.

- Roberts, K. (2015). "Dynamic Voting in Clubs". *Research in Economics*, 69(3):320–335.
- Rosén, A. (1997). "An Equilibrium Search-Matching Model of Discrimination". *European Economic Review*, 41(8):1589–1613.
- Schmeiser, S. (2012). "Corporate Board Dynamics: Directors Voting for Directors". *Journal of Economic Behavior & Organization*, 82(2):505 – 524.
- Selznick, P. (1948). "Foundations of the Theory of Organization". *American Sociological Review*, 13(1):25–35.
- Selznick, P. (1949). *TVA and the Grass Roots: A Study in the Sociology of Formal Organization*. Berkely: University of California Press.
- Tiebout, C. M. (1956). "A Pure Theory of Local Expenditures". *Journal of Political Economy*, 64:416–424.
- Zinov'yeva, N. and Bagues, M. (2015). "The Role of Connections in Academic Promotions". *American Economic Journal: Applied Economics*, 7(2):264–92.