Grant Lottery*

Why Funding Agencies May Rationally Introduce Randomness

Duk Gyoo Kim†

July 8, 2026

Abstract

Grant allocation processes occasionally exhibit seemingly random outcomes: high-quality proposals are often rejected and weaker ones are sometimes funded. Existing explanations attribute such randomness to reviewer error, inconsistency, or institutional inattention. This paper presents a different interpretation. Even when project quality is perfectly observed, a funding agency may *optimally* introduce randomness into the allocation mechanism. Randomness broadens the applicant pool at the cost of average quality, and an agency that values breadth alongside quality may rationally accept it, particularly when high-quality projects possess positive outside options. I develop a simple equilibrium model in which the agency chooses a degree of priority for high-quality proposals, and applicants decide strategically whether to apply. Agency welfare is hump-shaped in the degree of randomness: pure lotteries raise adverse selection because high-quality applicants with valuable outside options opt out, while strict meritocracy maximizes quality but excludes the broad base of low-quality applicants. When the agency values participation sufficiently, the optimal mechanism involves an interior level of randomness, rationalizing partial lotteries, random tie-breaking, and other randomness observed in grant allocation.

Keywords: grant allocation, lotteries, contest design, endogenous participation, research funding.

JEL Classification: C72, D47, I23, O38.

*I thank Jihwan Do, participants of Yonsei–Wuhan Microeconomics Workshop, and two anonymous reviewers for their constructive comments. This work was supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2023S1A5A8080738).

†School of Economics, Yonsei University, South Korea. Email: kim.dukgyoo@yonsei.ac.kr

1 Introduction

Resource allocation mechanisms often exhibit a puzzling degree of randomness. While we typically associate "lotteries" with gambling or revenue generation, a broader class of allocation mechanisms, ranging from military conscription to research funding, effectively functions as lotteries with participation thresholds. For instance, the selection of KATUSA (Korean Augmentation to the U.S. Army) soldiers requires applicants to meet a standardized English score, after which selection is purely random. Similarly, housing subscription systems often prioritize applicants based on meritocratic criteria (e.g., family size, savings duration) but revert to a lottery when the number of qualified applicants exceeds supply.

In the context of research funding, grant applicants often complain that application outcomes appear random: promising proposals are often rejected while weaker ones are sometimes funded. Despite continual refinements to peer review, such randomness persists across agencies, funding rounds, and scientific domains. This phenomenon is typically interpreted as a failure of evaluation: Reviewer disagreement, insufficient attention, and idiosyncratic criteria are widely believed to generate noise in final decisions. Accordingly, much of the literature focuses on how to reduce this noise to achieve an allocative efficiency in a meritocratic sense.

I argue that randomness may be an optimal feature of institutional design, rather than a flaw to be eliminated or merely accepted[1]. Even in an environment where reviewers perfectly observe project quality and the agency has no informational constraints, it may be optimal for the grant allocation mechanism to contain randomness. The central insight is that a grant competition is not simply a ranking exercise but a *strategic participation game*. Applicants enter only if they expect their probability of receiving a grant to justify the cost of applying relative to their outside option. Grant writing is itself a substantial investment[2], so the decision to apply is a genuine economic choice rather than a foregone conclusion. High-quality researchers typically possess stronger outside options: prior funding, institutional support, or better prospects for securing alternative grants. When the process becomes too random, these applicants lose their relative advantage and turn to their outside options, leaving an applicant pool dominated by low-quality projects. Strict meritocracy avoids this adverse selection and maximizes the quality of funded work, but it screens out the large

[1] Some studies claim that grant selection processes should actively consider randomness (Fang and Wagner, 2016; Avin, 2019; Philipps, 2022; Larue, 2026), but such a claim regards randomness as an unavoidable hurdle, rather than an optimal feature.

[2] von Hippel and von Hippel (2015) estimate that a single proposal absorbs on the order of a hundred principal-investigator hours.

base of low-quality applicants before they can enter, narrowing participation. A funding agency that values both the *quality* of funded projects and the *breadth of participation* thus faces a fundamental allocation tradeoff.

Participation matters for several reasons. The ultimate success of funded research is realized long after project selection, making ex post evaluation of program quality difficult or noisy. As a result, agencies rely on input-based indicators, in particular, the number and diversity of applicants, to assess the competitiveness, visibility, and inclusiveness of a program. Participation also broadens the space of ideas presented to reviewers, expanding the potential scope of funded research. For agencies with public or political accountability, broad participation signals openness and fairness.

In this paper, I construct a simple model capturing these considerations. There are high-quality (H) and low-quality (L) projects, with high-quality projects having a strictly positive outside option. The agency allocates a fixed number of grants G and chooses a *priority parameter* $\delta \in [0,1]$: a fraction δ of grants is awarded to H-types before any randomization occurs; the remainder are allocated randomly among remaining applicants. This model parsimoniously captures all selection mechanisms ranging from a perfect meritocracy ($\delta = 1$) to pure lottery ($\delta = 0$). Applicants decide whether to apply based on expected payoffs. Although the agency is modeled as selecting δ under perfect observation, δ is best read as the *effective priority* of the allocation, in practice determined jointly by deliberate design and by the accuracy of review and how strictly the ranking binds. A nominally meritocratic process run through noisy review delivers a lower effective δ than its rules suggest.

Even without any evaluation error, a pure lottery ($\delta = 0$) eliminates H's advantage and generates adverse selection, as high-quality applicants opt out. A fully meritocratic mechanism ($\delta = 1$) maximizes the quality of funded projects but excludes the broad base of low-quality applicants. I show that welfare, a weighted combination of funded-project quality and participation, is hump-shaped in δ, implying an interior optimum $\delta^* \in (0,1)$ under some conditions.

These results offer a *rationalization* for the randomness observed in grant allocation. The claim is not that funding agencies consciously implement lotteries—outside the explicit schemes noted above, there is little direct evidence that they do—but that randomness, whether deliberately built in or arising as a by-product of imperfect review, need not be a deviation from ideal review and may reflect, or unintentionally approximate, optimal institutional design. Section 6 makes this bridge precise: mild reviewer error is an imperfect substitute for deliberate randomization, so an agency whose effective priority already sits near the optimum has no reason to invest in eliminating it.

A pair of comparative statics sharpens the message: high-quality participation responds to priority in inverse proportion to the strength of applicants' outside options, so the strongest applicants whom the agency most wants to attract are the least drawn in by a more meritocratic rule. When randomization is optimal at all, the optimal degree of priority is then the highest level still consistent with full participation by the weaker pool, beyond which tightening the criteria trades breadth for quality one-for-one.

2 Related Literature

A growing body of empirical literature documents the use of randomness in high-stakes allocation. Blau et al. (2010) analyze the Committee on the Status of Women in the Economics Profession mentoring program, which used random selection to allocate presentation slots among junior female economists due to capacity constraints. In development economics, McKenzie (2017) studies the "YouWiN" competition in Nigeria, where business grants were awarded to a subset of qualified entrepreneurs via random selection, effectively creating a "threshold plus lottery" mechanism. Broader discussions, such as Frank (2016), argue that success in many competitive fields is heavily influenced by luck, and Pluchino et al. (2018) suggest that acknowledging this randomness can sometimes lead to more efficient outcomes than pretending to have perfect meritocratic precision. The current paper formalizes this intuition, showing exactly why an agency might strategically choose to retain randomness.

Another contribution of this study to the literature is to offer a different angle of viewing randomness in selection processes. Existing literature largely treats randomness as "error." A range of studies document low inter-rater reliability and the limited predictive validity of review scores: Cole et al. (1981) is the classic demonstration of the role of chance in peer review; Pier et al. (2018) and Bendiscioli (2019) document low agreement among reviewers of the same proposals; Graves et al. (2011) find that most funded grants would sometimes have gone unfunded once the random variability in panel scores is taken into account; and Fang et al. (2016) show that fine-grained percentile rankings are only weakly related to subsequent productivity. These findings motivate calls to reduce noise. In contrast, historical analyses of civil service reform, such as Moreira and Pérez (2024) on the Pendleton Act, highlight the trade-offs between discretion (which can be arbitrary) and rigid testing (which can be inefficient). I contribute to this debate by showing that even if "perfect" evaluation were possible, a welfare-maximizing agency might still prefer a mechanism that mimics the noise of imperfect review.

The model considered in this paper connects to contest theory literature involving outside options and participation costs. While previous work has explored lotteries for public goods provision (Morgan, 2000), vaccination incentives (Kim, 2021), or voting incentives (Gerardi et al., 2016), I focus specifically on the grant allocation problem. Most closely related, Gross and Bergstrom (2019) model grant competitions as contests and show that selecting at random from among proposals that clear a quality threshold can curb the effort researchers waste in preparing applications. Their rationale for randomization is the dissipation of proposal-writing effort. My rationale in this paper is distinct and complementary: partial randomization broadens participation when high-quality applicants hold outside options. Moreover, I demonstrate the optimal level of randomness even when proposals are costless to evaluate and reviewers are perfectly accurate. Unlike standard contest models where the goal is often to maximize aggregate effort, a typical funding agency faces a dual objective: securing high-quality projects while maintaining broad participation to ensure the program's visibility, inclusiveness, and legitimacy.

3 A Model

The applicant pool is a continuum of mass N[3], partitioned into a mass N_H of high-quality (H) projects and a mass N_L of low-quality (L) projects, with $N_H + N_L = N$. The agency awards a mass G of grants whose value is $V > 0$ and chooses the *priority granted to high-quality applicants*,[4] $\delta \in [0,1]$: a higher δ tilts the allocation toward H-types and a lower δ toward randomization, with $\delta = 1$ a fully meritocratic rule and $\delta = 0$ a pure lottery. Applicants who know their type and δ choose whether to apply. The application cost $c > 0$ is assumed to be the same for every applicant, and the value of the outside option varies, $O_H > O_L = 0$. Equivalently, one could let the two types share a common outside option but assign high-quality applicants a higher application cost; the two formulations are interchangeable. To avoid trivial cases where resources are too abundant, we assume $G < N_H$ so that competition among high-quality projects is non-trivial. Because the population is a continuum, δG is

[3]It would be more reasonable to assume the number of applicants is countable, but discretizations of the model do not qualitatively alter the main claims.

[4]Throughout the paper, *priority* is the primary term for δ. "Priority," "meritocracy," and the inverse of "randomness" are not synonyms in general—the first names the allocation rule, the latter two the character of its outcome—but in this model they are tied to δ monotonically: a higher δ means more priority for high-quality applicants, a more meritocratic rule, and less randomization. Where it reads more naturally, I therefore describe δ in terms of meritocracy or randomness, with no ambiguity intended.

also a mass and the priority parameter ranges freely over $[0,1]$.

Since we assume that quality of a project is perfectly observable upon application, the grants are allocated in the following way. Given δ, the agency awards $\min\{\delta G, A_H\}$ grants to H applicants, where A_H is the mass of H-type applications. If $\delta G < A_H$, the δG grants are allocated uniformly at random among the A_H applicants; if $\delta G \geq A_H$, all H applicants receive priority slots. The remaining $G - \min\{\delta G, A_H\}$ grants are allocated uniformly at random among the remaining applicants (remaining H applicants if any and all L applicants). Note that the priority parameter δ can be interpreted as the inverse of how much randomness plays a role in the allocation process. A perfect meritocracy is succinctly captured by $\delta = 1$: The grants are entirely allocated to H applicants if $A_H > G$, and L applicants may receive grants only when $A_H < G$. A perfect random selection is captured by $\delta = 0$: The entire $A_H + A_L$ applicants are equally likely to be awarded with probability $\frac{G}{A_H + A_L}$.

I focus on symmetric equilibria, in which all applicants of a given type use the same entry rule.[5] Let $p_H, p_L \in [0,1]$ denote the fractions of the high- and low-quality subpopulations that apply, so the applicant masses are $A_H = N_H p_H$ and $A_L = N_L p_L$. With a continuum of agents these masses are deterministic, which makes the per-applicant grant probabilities below exact rather than large-population approximations.

Writing $\rho(\delta, A_H, A_L) = \min\left\{1, \frac{G - \delta G}{A_H - \delta G + A_L}\right\}$ for the residual random-round probability, the per-applicant grant probabilities for (A_H, A_L) and δ are:

$$\pi_H(\delta, A_H, A_L) = \begin{cases} 1 & \text{if } A_H \leq \delta G, \\ \dfrac{\delta G}{A_H} + \dfrac{A_H - \delta G}{A_H}\rho & \text{if } A_H > \delta G, \end{cases}$$

$$\pi_L(\delta, A_H, A_L) = \begin{cases} \min\left\{1, \dfrac{G - A_H}{A_L}\right\} & \text{if } A_H \leq \delta G, \\ \rho & \text{if } A_H > \delta G. \end{cases}$$

These π_H and π_L follow from our deterministic priority slots (δG) among H followed by uniform randomization.[6]

[5] In the continuum, the coordination multiplicity that arises under perfect meritocracy with finitely many applicants (any combination of G applicants of the N_H high types applying and the rest staying out) collapses, and the symmetric strategy profile would be the natural selection.

[6] The min operators in π_L and ρ simply cap each probability at one. They bind only in the corner where applicants are scarcer than the grants still to be awarded, and we set $\pi_L = 0$ when $A_L = 0$. Neither cap binds on the equilibrium path: in particular, $A_H < \delta G$ never arises in equilibrium, because in that case, every high type is funded with certainty ($\pi_H = 1 > \tau_H$), so high types keep applying until $A_H \geq \delta G$.

Now we turn to the perspective of each applicant. Type i applies if and only if:

$$\pi_i(\delta, A_H, A_L)V - c \geq O_i,$$

with $O_H > 0$, $O_L = 0$. It is convenient to write the two participation thresholds as $\tau_H = \frac{c+O_H}{V}$ and $\tau_L = \frac{c}{V}$, with $0 < \tau_L < \tau_H$; type i applies whenever its grant probability is at least the corresponding threshold probability τ_i.

We maintain three assumptions throughout.

Assumption 1. $\tau_H < 1$*, that is,* $V > c + O_H$.

Assumption 1 postulates that a type-H applicant can profitably apply. Otherwise, no high-quality applicant would apply regardless of the allocation rule.

Assumption 2. $N_H \geq G/\tau_H$

Assumption 2 postulates that the meritocratic mass G/τ_H is feasible. This assumption implies that the highest average quality of accepted applications is achievable, making the benchmark for the maximized quality straightforward. This assumption strengthens the model setup $G < N_H$, since $G/\tau_H > G$.

Assumption 3. $0 < N_L < G/\tau_L$

Assumption 3 states that the entire low-quality pool applies when priority is low ($\delta = 0$) and the expected lottery margin is large enough to compensate for the application cost.

Agency welfare is defined as follows:

$$W(\delta) = \alpha Q(\delta) + (1-\alpha)A(\delta),$$

where

$$Q(\delta) = \frac{\text{expected mass of H winners}}{G}, \qquad A(\delta) = \frac{A_H(\delta) + A_L(\delta)}{N},$$

and $\alpha \in [0,1]$. Normalizing the participation rate by N keeps both components in $[0,1]$, so that W is a convex combination of funded quality and breadth of participation. It is important to note that we read W as the agency's objective rather than as social welfare: the weight on the participation breadth, $1-\alpha$, captures the agency's concern for program visibility, legitimacy, perceived accessibility and accountability, and quantitative metrics such as application or submission volume and the demographic diversity of the pool.

4 Equilibrium Analyses

Applicants simultaneously choose whether to apply. Write $U_i(\delta) = \pi_i(\delta, A_H, A_L)V - c - O_i$ for the net payoff a type-i applicant obtains from applying. A symmetric equilibrium is a pair $(p_H(\delta), p_L(\delta))$, with $A_H = N_H p_H$ and $A_L = N_L p_L$, satisfying for each $i \in \{H, L\}$ the complementarity conditions

$$p_i = 0 \Rightarrow U_i \leq 0, \qquad p_i \in (0,1) \Rightarrow U_i = 0, \qquad p_i = 1 \Rightarrow U_i \geq 0.$$

Proposition 1 (Equilibrium Existence and Uniqueness)**.** *Under Assumptions 1–3, for any* $\delta \in [0,1]$ *there exists a unique symmetric equilibrium* $(p_H(\delta), p_L(\delta))$*.*

Proof. See Appendix A. □

We first explore some equilibrium properties when δ is extreme at 0 and 1, and then characterize the equilibrium application masses for general δ.

Under $\delta = 1$, all grants are reserved for H types, and since $N_H \geq G$ the high-quality applicants compete among themselves, each winning with probability G/A_H. In symmetric equilibrium, they enter until this probability equals their threshold $\frac{c+O_H}{V}$, so $A_H = \frac{GV}{c+O_H}$. This is the largest high-quality participation attainable under any δ as greater priority never deters H's entry. Every funded project is high-quality, so the quality component Q is maximal. Low-quality types, however, never apply, so the applicant pool is at its narrowest. Strict meritocracy thus secures quality at the cost of participation breadth.

Meanwhile, pure lotteries generate adverse selection. Under $\delta = 0$, there is no priority for H, so every applicant wins with the common probability $\rho = \min\{1, \frac{G}{A_H + A_L}\}$: the two types face identical chances, but the high type requires the larger return ($\tau_H > \tau_L$) and so is the first to drop out as the pool fills. When the low-quality pool can fill the lottery on its own, that is, when $N_L \geq \frac{G}{\tau_H}$, high types opt out entirely, implying that $A_H = 0$, $A_L = \min\{N_L, \frac{G}{\tau_L}\}$, and every funded project is low-quality, $Q(0) = 0$. When instead $N_L < \frac{G}{\tau_H}$, the low-quality pool does not fill all the grants, so some high types remain, $A_H = \frac{G}{\tau_H} - N_L$, but funded quality $Q(0) = 1 - N_L \frac{\tau_H}{G}$ is still strictly below the meritocratic level. Either way, removing priority lowers the quality of funded work—sharply so when the low-quality mass is large.

The equilibrium admits a closed form. Recall $\tau_H = \frac{c+O_H}{V}$ and $\tau_L = \frac{c}{V}$. As δ rises, it passes through three regimes, which are described in Proposition 2.

Proposition 2 (Equilibrium application masses)**.** *Under Assumptions 1–3, the meritocratic mass* $\frac{G}{\tau_H}$ *is feasible and the entire low-quality pool applies at low* δ*. Let* $\widehat{A}_H(\delta)$ *be the unique* *root in* $(\delta G, \infty)$ *of* $\delta G + \frac{(A_H - \delta G)(G - \delta G)}{A_H - \delta G + N_L} = \tau_H A_H$. *The equilibrium mass of high-quality applicants is*

$$A_H^*(\delta) = \begin{cases} \widehat{A}_H(\delta), & 0 < \delta \le \delta_1, \\ \dfrac{\delta G(1-\tau_L)}{\tau_H - \tau_L}, & \delta_1 \le \delta \le \delta_2, \\ \dfrac{G}{\tau_H}, & \delta_2 \le \delta \le 1, \end{cases}$$

where $\delta_2 = \frac{\tau_H - \tau_L}{\tau_H(1-\tau_L)}$ *and* δ_1 *is the unique* δ *at which* $\pi_L(\delta, A_H^*, N_L) = \tau_L$*. In the first regime* $(\delta \in (0, \delta_1))$*, all low-quality applicants apply* $(A_L = N_L)$*; in the second* $(\delta \in [\delta_1, \delta_2])$*, both types mix and* $A_L^*(\delta) = \frac{G - \delta G}{\tau_L} - A_H^*(\delta) + \delta G$ *is strictly decreasing in* δ*; in the third regime* $(\delta \in [\delta_2, 1])$*, low-quality applicants are excluded* $(A_L = 0)$ *and every grant goes to an* H *type, so* $Q(\delta) = 1$.

Proof. See Appendix A. □

Two consequences are worth noting. First, on the interior regime $[\delta_1, \delta_2]$, high-quality participation $A_H^*(\delta) = \delta \frac{G(1-\tau_L)}{\tau_H - \tau_L}$ is exactly proportional to δ, so wherever the boundaries δ_1 and δ_2 happen to fall, the high-type response to priority δ is linear within that region, a structural property of the equilibrium rather than an artifact of particular parameters. Second, δ_1 is the largest priority at which the *entire* low-quality pool still participates; beyond it, raising priority buys quality only by shrinking the applicant pool.

The two components of agency welfare, $Q(\delta)$ and $A(\delta)$, pull in opposite directions. Funded quality $Q(\delta)$ is non-decreasing in δ, since priority monotonically steers grants toward high-quality projects. Meanwhile, the participation rate $A(\delta)$ is single-peaked, as shown in the following lemma.

Lemma 1. *Fix an equilibrium with* $A_H^* > \delta G$*. Then* $A_H^*(\delta)$ *is nondecreasing in* δ*, with* A_H^* *strictly increasing for small* δ *and locally constant at* $A_H^* = \frac{GV}{c + O_H}$ *near* $\delta = 1$*, while* $A_L^*(\delta)$ *is nonincreasing in* δ *and falls to zero as* $\delta \to 1$*. Consequently, the participation rate* $A(\delta) = \frac{A_H^* + A_L^*}{N}$ *is single-peaked: it rises near* $\delta = 0$ *and falls near* $\delta = 1$.

Proof. See Appendix A. □

Lemma 1 states that as δ rises from zero, high-quality applicants enter and adverse selection abates, but beyond a point the low-quality pool is progressively screened out, so participation eventually falls. Thus, agency welfare $W = \alpha Q + (1-\alpha)A$ trades off the quality component pulling the optimum toward full meritocracy with the participation component toward the interior peak, so the optimal priority depends on how much the agency values breadth.

Proposition 3 (Interior optimal randomness)**.** *There is a threshold* $\bar{\alpha} \in (0,1)$ *such that: if* $\alpha < \bar{\alpha}$*, the unique maximizer is the interior kink* $\delta^* = \delta_1$*; if* $\alpha > \bar{\alpha}$*, every priority on the meritocracy plateau* $[\delta_2, 1]$ *is optimal; and at* $\alpha = \bar{\alpha}$ *the agency is indifferent across all* $\delta \in [\delta_1, 1]$.

Proof. See Appendix A. □

Proposition 3 is the paper's central result: an agency that places enough weight on breadth optimally chooses an interior degree of priority, neither the pure lottery ($\delta = 0$) which invites adverse selection nor the full meritocracy ($\delta = 1$) which narrows the applicant pool. It thereby rationalizes the threshold-plus-lottery mechanisms highlighted in the introduction, e.g., KATUSA selection above an English-score cutoff, housing subscriptions that randomize among qualified applicants, and the partial lotteries increasingly used in grant review, as optimal institutional design rather than administrative expedience. The driving force is a participation effect: because quality is assumed to be perfectly observed, the optimal randomness described in Proposition 3 reflects a deliberate trade of quality for breadth, not a concession to noisy evaluation. The threshold $\bar{\alpha}$ then separates the agencies that should build in randomness from those that should select purely on merit.

5 A Numerical Illustration

The equilibrium and the optimal priority are available in closed form, so the following worked example is illustrative rather than evidentiary: it does not establish the results—these are analytical and hold for all admissible parameters—but it makes their content concrete. Because the closed-form masses are piecewise and depend jointly on several primitives, a single calibration is the most transparent way to see how the three regimes fit together, how the quality and participation components trade off as priority varies, and where the interior optimum falls. A worked example also fixes magnitudes that the algebra leaves implicit, conveying not just the shape of the welfare function but how pronounced its interior peak is.

The figures below plot the equilibrium objects for a representative parameterization, $N_H = 120$, $N_L = 100$, $G = 30$, $V = 10$, $c = 1$, and $O_H = 2$, so that the participation thresholds are $\tau_H = \frac{c+O_H}{V} = 0.3$ and $\tau_L = \frac{c}{V} = 0.1$. The curves are computed from the equilibrium conditions of Section 3.

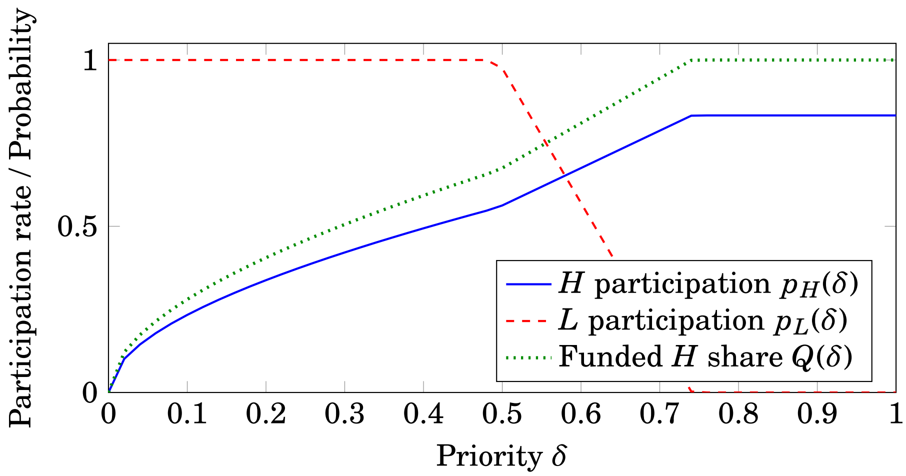

Figure 1: Equilibrium participation rates p_H, p_L, and the funded high-quality share Q.

Figure 1 traces the equilibrium as priority δ varies. High-quality participation p_H is lowest under the pure lottery, as adverse selection keeps strong applicants out as they pursue their outside options, and it rises with δ as priority restores their advantage, following the straight line of the interior regime (Proposition 2) before flattening once meritocracy caps high-type entry at $A_H = \frac{G}{\tau_H}$. Low-quality participation p_L moves in the opposite direction: the entire low-type mass applies until $\delta_1 \approx 0.49$, after which rising priority rations it out, reaching zero at $\delta_2 \approx 0.74$. The funded high-quality share Q climbs steadily to one as the pool tilts toward H types. Breadth is therefore widest at low-to-intermediate priority while quality peaks at high priority, which demonstrates the tension the agency must resolve.

Figure 2 decomposes the agency's objective for $\alpha = 0.4$. The quality component Q rises monotonically and reaches one once low types are excluded for $\delta \geq \delta_2$, while the participation component A is single-peaked (Lemma 1), attaining its maximum at δ_1, where the high-type inflow just offsets the onset of low-type exclusion. Their weighted sum $W(\delta) = \alpha Q(\delta) + (1-\alpha)A(\delta)$ depicted in solid line is hump-shaped, and because $0.4 = \alpha < \bar{\alpha} \approx 0.48$ it peaks at the interior kink $\delta^* = \delta_1 \approx 0.49$, which implies that the agency wants to implement a partial lottery (Proposition 3). Raising α above $\bar{\alpha}$ would move the optimum onto the meritocratic plateau.

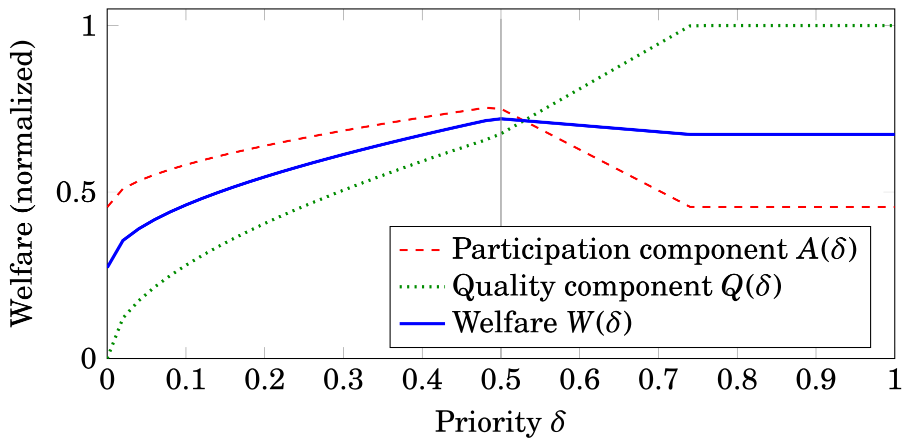

Figure 2: Decomposition of the agency's objective $W = \alpha Q + (1-\alpha)A$ for $\alpha = 0.4$ (same parameters as Figure 1).

6 Discussion and Extensions

6.1 The Cost of Perfect Evaluation

Our baseline model assumes that the agency can observe project quality without cost, yet we still find that the optimal mechanism involves randomness ($\delta^* < 1$). In practice, evaluation is costly. As noted in the context of civil service exams or comprehensive grant reviews, achieving a strictly meritocratic ranking requires significant investment in reviewer time and expertise. If we were to introduce a marginal cost of evaluation r into the model, the agency would have a secondary, budgetary incentive to reduce the precision of the ranking and rely more on randomization. A parallel cost arises on the applicant side: Gross and Bergstrom (2019) show that competition dissipates proposal-writing effort, which randomization above a threshold can curb. Thus, the breadth-versus-quality tradeoff we identify acts as a complement to these standard cost-saving rationales for lotteries. Our result suggests that agencies should not strive for perfect accuracy at all costs; rather, a "good enough" evaluation combined with a lottery may be both strategically and fiscally optimal.

The same logic extends, even more directly, to the *accuracy* of review rather than its cost. In the baseline model, the agency chooses the degree of priority directly. In practice, however, review precision and effective priority are often administratively bundled: making review more accurate usually means making the final ranking more binding. In such an environment, a costless improvement in review accuracy need not be desirable. If perfect accuracy would effectively implement strict meritocracy, corresponding to $\delta = 1$, then an agency whose optimum lies at an interior δ would rationally reject even a free reform that eliminates all review noise. The reason is not that information is intrinsically harmful, but that the institutional use of that information may move the mechanism beyond the agency's optimal degree of priority.

This interpretation reframes reviewer error. Noise is not the fundamental justification for randomness; the participation channel is. Nevertheless, mild reviewer error can serve as an imperfect substitute for deliberate randomization by lowering the effective priority of high-ranked proposals. Because deliberate randomization and reviewer noise both act only through the effective priority, two mechanisms that induce the same effective priority deliver the same allocation; at that level the randomness an agency builds in and the randomness it inherits are indistinguishable. Such noise is valuable only up to a point. If review becomes too noisy, the mechanism approaches the pure-lottery limit, high-quality applicants with outside options opt out, and adverse selection reappears. The model therefore implies an interior optimal precision of review mirroring the interior optimal priority: accurate enough to avoid adverse selection, but not so accurate, or not so binding, that it forecloses the participation benefits of partial randomness.

6.2 Behavioral Implications

The current model assumes applicants rationally calculate their winning probabilities π_H and π_L. However, behavioral economics suggests that agents, particularly those with lower outside options, may overestimate their chances in lottery-like systems. If low-quality applicants perceive the "random" component of the allocation as providing a "false hope" of success, their participation may be higher than predicted by standard theory. From the agency's perspective, this behavioral bias could further justify lower values of δ, as it allows the agency to maintain high application volume or participation metrics without significantly diluting the success rate of high-quality projects.

This bias may do more than raise participation; it may make participation cheap to buy. Because over-optimistic low types overweight a small winning probability, even a sliver of randomness, a δ slightly smaller than one, can draw them in disproportionately, so the agency secures a large gain in breadth while diverting only a handful of grants from the meritocratic ranking. The visible possibility that anyone might "get lucky" does the work, which suggests a testable pattern: selection processes that preserve a hint of unpredictability may attract more applications, and more application revenue (if application requires fees) than otherwise identical predictable ones, even when the true odds are almost unchanged.

6.3 Policy Implications

Our framework serves as a diagnostic tool for evaluating current allocation mechanisms. For example, highly over-subscribed systems like housing lotteries or popular research grants often operate near the lottery end of the spectrum ($\delta \approx 0$). Our model suggests that if high-quality applicants (e.g., families with urgent needs or researchers with breakthrough ideas) are opting out due to the sheer noise of the process, policymakers should consider increasing δ—essentially tightening the "qualification threshold" before the lottery kicks in. Conversely, strictly meritocratic systems that suffer from low participation diversity might benefit from introducing explicit randomization layers, similar to the "Explorer Grants" or the "YouWiN" model. More broadly, the model is not confined to threshold lotteries that randomize only among proposals already judged fundable. By letting the lower-priority pool retain a positive chance, it captures any environment in which reducing the priority of top-ranked applications raises the perceived opportunity of weaker applicants and thereby shapes who enters; the threshold-only lottery is the special case of large δ.

The model's prescription depends on α, and the comparative statics give that weight a direct institutional reading. Where α is high, so that the marginal value of funding the single strongest project dominates, near-meritocracy is optimal. Frontier-research funders, whose mandate is to support the most promising science, are the natural example. Where α is low, the principal values breadth enough that an interior degree of randomness becomes optimal, and such environments are common: broad-access subsidy and policy-uptake programs are often judged on take-up and reach; public schemes are often held accountable for the number and demographic diversity of the applicants they attract; and editorial or admissions processes treat submission volume and selectivity as indicators of their standing. A single agency may even comprise people with different viewpoints on the quality weight α, applying sharp priority to its flagship competitions while deliberately broadening entry where participation is itself part of a program's value. Thus, the main policy implication might not be the mechanical introduction of randomness, but the importance of identifying the agency's objective α before choosing the optimal degree of randomness.

7 Concluding Remarks

This paper offers a strategic rationale for randomness in grant allocation that does not rest on reviewer error. When an agency values the breadth of participation alongside the quality of funded work, and high-quality applicants hold outside options, neither extreme is optimal: a pure lottery induces adverse selection, while strict meritocracy screens out the broad base of applicants whose participation the agency values. An interior degree of priority—partial randomization—maximizes welfare whenever participation receives sufficient weight, rationalizing partial lotteries and random tie-breaking as features of optimal design rather than failures of evaluation.

The model is deliberately stripped down, and its simplifications mark natural extensions. Quality is binary and costlessly observed; applicant effort is exogenous; and the agency's valuation of participation is taken as a primitive rather than derived from an explicit accountability or option-value problem. Section 6 indicates how costly evaluation and behavioral misperception would reinforce, rather than overturn, the interior-randomness result. Endogenizing effort and quality, and grounding the participation objective in a richer institutional model, are left for future work.

Perhaps the most natural extension is to make the model dynamic. Grant competitions recur, and applicants do not observe the priority parameter δ directly; they infer it from the public record of past rounds about which proposals were funded and which were passed over, and update their beliefs before deciding whether to apply again. Where the static model takes applicants' beliefs as given, a dynamic version would let them form from experience, turning δ into an object that applicants learn about and making participation depend on the history of realized allocations rather than on a commonly known rule. This raises questions the static analysis cannot reach: whether the interior optimum survives when the agency's randomness must be *inferred* rather than announced; whether learning sustains multiple long-run regimes, from a broad-participation equilibrium held up by optimistic beliefs to a collapsed equilibrium in which adverse selection dominates; and whether an agency can shape the observable record of winners to manage those beliefs. This dynamic, learning-based extension would be a central direction for future work.

References

Avin, Shahar, "Mavericks and lotteries," *Studies in History and Philosophy of Science Part A*, 2019, *76*, 13–23.

Bendiscioli, Sandra, "The troubles with peer review for allocating research funding," *EMBO reports*, 2019, *20* (12), e49472.

Blau, Francine D., Janet M. Currie, Rachel T. A. Croson, and Donna K. Ginther, "Can Mentoring Help Female Assistant Professors? Interim Results from a Randomized Trial," *American Economic Review*, May 2010, *100* (2), 348–352.

Cole, Stephen, Jonathan R. Cole, and Gary A. Simon, "Chance and consensus in peer review," *Science*, 1981, *214* (4523), 881–886.

Fang, Ferric C. and Daniel M. Wagner, "Research Funding: the Case for a Modified Lottery," *mBio*, 2016, *7* (2), e00422–16.

_, Anthony Bowen, and Arturo Casadevall, "NIH peer review percentile scores are poorly predictive of grant productivity," *eLife*, 2016, *5*, e13323.

Frank, Robert H., *Success and Luck: Good Fortune and the Myth of Meritocracy*, Princeton, NJ: Princeton University Press, 2016.

Gerardi, Dino, Margaret A. McConnell, Julian Romeo, and Leeat Yariv, "Get Out the Costly Vote: Institutional Design for Greater Participation," *Economic Inquiry*, 10 2016, *54* (4), 1963–1979.

Graves, Nicholas, Adrian G. Barnett, and Philip Clarke, "Funding grant proposals for scientific research: retrospective analysis of scores by members of grant review panel," *BMJ*, 2011, *343*, d4797.

Gross, Kevin and Carl T. Bergstrom, "Contest models highlight inherent inefficiencies of scientific funding competitions," *PLoS Biology*, 2019, *17* (1), e3000065.

Kim, Duk Gyoo, "Vaccination Lottery," *Economics Letters*, 2021, *208*, 110059.

Larue, Louis, "Funding Research Randomly," *Journal of Applied Philosophy*, 2026, *43* (1), 258–275.

McKenzie, David, "Identifying and Spurring High-Growth Entrepreneurship: Experimental Evidence from a Business Plan Competition," *American Economic Review*, August 2017, *107* (8), 2278–2307.

Moreira, Diana and Santiago Pérez, "Civil Service Exams and Organizational Performance: Evidence from the Pendleton Act," *American Economic Journal: Applied Economics*, 7 2024, *16* (3), 250–291.

Morgan, John, "Financing Public Goods by Means of Lotteries," *The Review of Economic Studies*, 2000, *67* (4), 761–784.

Philipps, Axel, "Research funding randomly allocated? A survey of scientists' views on peer review and lottery," *Science and Public Policy*, 06 2022, *49* (3), 365–377.

Pier, Elizabeth L., Markus Brauer, Amarette Filut, Anna Kaatz, Joshua Raclaw, Mitchell J. Nathan, Cecilia E. Ford, and Molly Carnes, "Low agreement among reviewers evaluating the same NIH grant applications," *Proceedings of the National Academy of Sciences*, 2018, *115* (12), 2952–2957.

Pluchino, Alessandro, Alessio E. Biondo, and Andrea Rapisarda, "Talent vs. Luck: The Role of Randomness in Success and Failure," *Advances in Complex Systems*, 2018, *21* (03n04), 1850014.

von Hippel, Ted and Courtney von Hippel, "To Apply or Not to Apply: A Survey Analysis of Grant Writing Costs and Benefits," *PLOS ONE*, 03 2015, *10* (3), 1–8.

A Proofs

A.1 Proof of Proposition 1

Proof. Fix δ. Recall Assumption 1, $\tau_H < 1$ and Assumption 2, $N_H \geq \frac{G}{\tau_H}$. If $A_H \leq \delta G$ then every high type is funded with certainty, so $\pi_H = 1 > \tau_H$ and $U_H > 0$; all high types then apply, giving $A_H = N_H > \delta G$ (as $N_H \geq G/\tau_H > \delta G$), a contradiction. Hence every equilibrium has $A_H > \delta G$, where, with $\rho = \min\{1, (G - \delta G)/(A_H - \delta G + A_L)\}$,

$$\pi_H = \frac{\delta G}{A_H} + \frac{A_H - \delta G}{A_H}\,\rho = \pi_L + \frac{\delta G}{A_H}(1 - \pi_L), \qquad \pi_L = \rho.$$

The low type's best response. For fixed $A_H > \delta G$, π_L is continuous and non-increasing in A_L, strictly so where $\pi_L < 1$. Hence the best response is $A_L = N_L$ if $\pi_L(A_H, N_L) \geq \tau_L$, $A_L = 0$ if $\pi_L(A_H, 0) \leq \tau_L$, and otherwise the unique $A_L \in (0, N_L)$ with $\pi_L = \tau_L$. Denote it $\ell(A_H)$; it is continuous, and the residual mass $M(A_H) = A_H - \delta G + \ell(A_H)$ is non-decreasing in A_H—it increases with A_H in either corner regime and equals the constant $(G - \delta G)/\tau_L$ where the low type is interior. Thus $\pi_L = \rho$ is non-increasing in A_H along ℓ.

The high type pins a unique A_H. Let $\varphi(A_H) = \pi_H\big(A_H, \ell(A_H)\big)$. Differentiating the second expression for π_H along ℓ,

$$\varphi'(A_H) = \frac{A_H - \delta G}{A_H}\frac{d\pi_L}{dA_H} - \frac{\delta G}{A_H^2}(1-\pi_L),$$

which is strictly negative wherever $\pi_L < 1$, since both terms are non-positive and the second is strictly negative. Hence φ is continuous, equals 1 for small A_H (where the residual pool is scarce and $\rho = 1$), and is strictly decreasing thereafter. Since $\varphi(A_H) \to 1 > \tau_H$ as $A_H \downarrow \delta G$ and $\tau_H < 1$, the equation $\varphi(A_H) = \tau_H$ has at most one root in $(\delta G, N_H)$: if $\varphi(N_H) \geq \tau_H$ the unique high-type best response is the corner $A_H = N_H$; otherwise φ crosses τ_H once, at a unique interior A_H. Either way A_H exists and is unique, and then $A_L = \ell(A_H)$ is unique. The symmetric equilibrium therefore exists and is unique. □

A.2 Proof of Proposition 2

By Proposition 1 the symmetric equilibrium is unique for each δ, and (as shown there) it has $A_H > \delta G$, so

$$\pi_L = \frac{G - \delta G}{A_H - \delta G + A_L}, \qquad \pi_H = \frac{\delta G}{A_H} + \frac{A_H - \delta G}{A_H}\pi_L$$

with the caps slack on the equilibrium path. We compute the equilibrium by solving the binding participation conditions, working from the interior regime outward.

Interior regime. Suppose both types mix, so $U_H = U_L = 0$, i.e., $\pi_H = \tau_H$ and $\pi_L = \tau_L$. From $\pi_L = \tau_L$,

$$A_H - \delta G + A_L = \frac{G - \delta G}{\tau_L}.$$

Substituting $\pi_L = \tau_L$ into $\pi_H = \tau_H$ gives $\frac{\delta G}{A_H} + \frac{A_H - \delta G}{A_H}\tau_L = \tau_H$, i.e., $\delta G(1-\tau_L) = (\tau_H - \tau_L)A_H$, so

$$A_H^*(\delta) = \frac{\delta G(1-\tau_L)}{\tau_H - \tau_L}, \qquad A_L^*(\delta) = \frac{G - \delta G}{\tau_L} - A_H^*(\delta) + \delta G.$$

Thus A_H^* is linear and increasing in δ while A_L^* is decreasing. This regime applies wherever $0 \leq A_L^* \leq N_L$. Setting $A_L^* = 0$ gives $A_H^* = G/\tau_H$ and the upper boundary $\delta_2 = (\tau_H - \tau_L)/[\tau_H(1-\tau_L)]$; the lower boundary δ_1 is the value at which $A_L^* = N_L$, equivalently $\pi_L(\delta, A_H^*, N_L) = \tau_L$. By Assumption 2, $0 < N_L < \frac{G}{\tau_L} = A_L^*(0)$; since A_L^* decreases continuously from $A_L^*(0) = \frac{G}{\tau_L}$ to $A_L^*(\delta_2) = 0$, it crosses N_L exactly once, so $0 < \delta_1 < \delta_2$ and both the bottom and interior regimes are non-degenerate.

Bottom regime $(0 < \delta \leq \delta_1)$. For $\delta < \delta_1$ the interior value of A_L^* exceeds N_L, so the low type is at the corner $A_L = N_L$ (all apply, $\pi_L \ge \tau_L$). The high type is interior, and A_H^* is the unique root of $\pi_H(\delta, A_H, N_L) = \tau_H$, namely

$$\delta G + \frac{(A_H - \delta G)(G - \delta G)}{A_H - \delta G + N_L} = \tau_H A_H.$$

Top regime ($\delta_2 \le \delta \le 1$). Here $A_L^* = 0$; with no low types, $\pi_H = G/A_H = \tau_H$ gives $A_H^* = G/\tau_H$, and every grant goes to a high type, so $Q(\delta) = 1$.

The masses are feasible throughout: $A_H^* \le G/\tau_H \le N_H$ by Assumption 2, $N_H \ge G/\tau_H$, and $0 \le A_L^* \le N_L$ by construction. Together with the uniqueness from Proposition 1, this establishes the proposition. □

A.3 Proof of Lemma 1

For $A_H > \delta G$, the high-quality probability is $\pi_H = \frac{\delta G}{A_H} + \frac{A_H - \delta G}{A_H} \cdot \frac{G - \delta G}{A_H - \delta G + A_L}$, which is strictly decreasing in A_H. Hence, the equilibrium mass of high-quality applicants A_H^* is pinned by the indifference condition $\pi_H(\delta, A_H^*, A_L^*)V = c + O_H$. Direct differentiation shows $\frac{\partial \pi_H}{\partial \delta} \ge 0$ at fixed (A_H, A_L): shifting a marginal slot from the shared random pool into a priority slot reserved for H raises H's probability of being funded. Since π_H is decreasing in A_H, restoring indifference after an increase in δ requires a (weakly) larger A_H^*, so A_H^* is nondecreasing in δ; the effect is strict near $\delta = 0$. Near $\delta = 1$, all priority slots already accrue to H, $\pi_H \to \frac{G}{A_H}$, and $A_H^* \to \frac{GV}{c+O_H}$ independently of δ, so A_H^* is locally constant. For L types, as δ rises the random pool shrinks ($G - \delta G$ falls) while the residual H mass competing in it rises, so π_L falls and A_L^* is nonincreasing, reaching zero at $\delta = 1$. Single-peakedness of A follows from A_H^* rising and A_L^* falling at rates that reverse the sign of $A'(\delta)$ between the boundaries. Continuity and the implicit-function steps are standard; Proposition 2 gives the exact masses, with $A(\delta)$ peaking at the kink δ_1. □

A.4 Proof of Proposition 3

By Proposition 2, the equilibrium masses are piecewise in δ: on the bottom regime $A_L = N_L$ and A_H^* rises; on the interior regime A_H^* rises linearly while A_L^* falls linearly; on the top regime $A_H^* = \frac{G}{\tau_H}$ and $A_L^* = 0$ are constant. Hence $Q(\delta) = A_H^*(\delta)\frac{\tau_H}{G}$ is nondecreasing and equals 1 on the top regime, while the participation rate $A(\delta) = \frac{A_H^* + A_L^*}{N}$ rises on the bottom regime, attains its maximum at the kink δ_1, and falls on the interior regime. Thus $W = \alpha Q + (1 - \alpha)A$ is continuous, increasing on the bottom regime, linear on the interior regime, and constant on the top regime.

On the interior regime Q and A have the constant slopes

$$Q'_B = \frac{\tau_H(1-\tau_L)}{\tau_H - \tau_L} > 0, \qquad A'_B = \frac{G(1-1/\tau_L)}{N} < 0,$$

so W there has slope $\alpha Q'_B + (1-\alpha)A'_B$, which is negative iff $\alpha < \bar{\alpha}$, where

$$\bar{\alpha} = \frac{-A'_B}{Q'_B - A'_B} \in (0,1).$$

For $\alpha < \bar{\alpha}$, W rises up to δ_1 and falls thereafter, so it is maximized at the kink $\delta^* = \delta_1$—the largest priority at which all low-quality applicants still apply. For $\alpha > \bar{\alpha}$, W rises across the interior regime and is constant on $[\delta_2, 1]$, so every point of the meritocracy plateau is optimal, in particular full meritocracy. At $\alpha = \bar{\alpha}$ the interior slope vanishes, so W is flat on $[\delta_1, 1]$ and the agency is indifferent across that range. No global differentiability of W is invoked; the interior maximum is a kink. □