Reformulating Release Figure Methodology: A Multi-Objective Approach to Equitable Chapter Outcomes in Structured Sorority Recruitment

Alexis Harvey
LSU A&M
Baton Rouge, Louisiana, USA
aharv25@lsu.edu

Abstract

Release Figure Methodology (RFM) is the allocation mechanism that governs structured sorority recruitment on campuses affiliated with the National Panhellenic Conference. RFM sets, for each round of recruitment, the number of potential new members (PNMs) each chapter may invite back, with the stated objective of maximizing the number of PNMs placed into a chapter. We argue that this single-objective formulation is well specified for aggregate placement but structurally silent on the distribution of outcomes across chapters, and that this silence produces a self-reinforcing divergence between high-demand and low-demand chapters. We formalize the current mechanism as a constrained optimization problem, identify four failure modes that follow from its objective function and its inputs, and propose a reformulation comprising six changes: an ε-constrained lexicographic objective that selects among near-optimal allocations by minimizing dispersion in chapter fill rate; retention weighting of projected placements; cardinal rather than purely ordinal elicitation of PNM preferences; temporal smoothing of release figures to interrupt year-over-year compounding; quota expressed as a band rather than a point value; and a total-setting rule anchored to the median chapter size with an automatic dispersion trigger. We specify an evaluation protocol based on historical backtesting and agent-based counterfactual simulation, define the metrics required to characterize the price of fairness of each change, and state falsifiable hypotheses. The contribution is a decision framework that makes the equity and efficiency tradeoff in recruitment allocation explicit and measurable rather than implicit and unexamined.

CCS Concepts

• **Theory of computation** → **Algorithmic game theory and mechanism design**; • **Information systems** → *Decision support systems*; • **Applied computing** → *Education*.

Keywords

matching markets, mechanism design, resource allocation, fairness, student organizations, higher education administration

1 Introduction

Structured sorority recruitment is a two-sided matching market operating under unusually tight constraints. Several hundred to several thousand potential new members (PNMs) and a fixed roster of chapters must be matched within a compressed window, typically under a week, through a sequence of mutual-selection rounds that narrow monotonically until each PNM holds preferences over at most two chapters. Unlike most matching markets studied in the mechanism design literature [5–7], the number of participants each side may carry forward into the next round is not chosen by the participants. It is set centrally, in advance of each round, by an allocation procedure known as Release Figure Methodology (RFM).

RFM determines each chapter's *release figure*: the number of PNMs the chapter may invite to the next round. Release figures are derived from observed demand, from quota (the number of new members each chapter may pledge), and from total (the maximum chapter size permitted on that campus). The mechanism's declared objective is to maximize the number of PNMs who receive and accept a bid. By that measure, RFM performs well. Campuses that adopted it report materially higher placement rates than the ad hoc practices that preceded it.[1]

Aggregate placement, however, is one coordinate of the outcome vector. The same number of total placements is consistent with a wide range of distributions across chapters, and the choice among those distributions is not neutral. A chapter that repeatedly finishes recruitment below quota loses membership dues, loses recruitment workforce, loses the visibility that drives the following year's demand, and in the terminal case loses its charter. The mechanism does not cause the underlying demand asymmetry. It does, however, decide how the consequences of that asymmetry are distributed, and it currently makes that decision as an unexamined byproduct of optimizing a different quantity.

This paper makes that decision explicit. Section 2 describes the institutional setting. Section 3 formalizes RFM as a constrained optimization problem. Section 4 identifies four structural failure modes. Section 5 proposes a reformulation. Section 6 specifies an evaluation protocol and states testable hypotheses. Section 7 addresses limitations and ethical constraints on deployment.

Scope. We take chapter demand as exogenous. Nothing proposed here makes a chapter more attractive to PNMs; that is a function of chapter programming, reputation, and recruitment execution, and lies outside the allocation mechanism. The claim of this paper is narrower and, we believe, defensible: an allocation mechanism should not amplify an existing asymmetry when a near-equivalent allocation that does not amplify it is available.

[1] NPC operational documents, including the *Manual of Information* and RFM specialist training materials, are distributed to member organizations and are not publicly archived. This statement is taken from public data and statistics published on an institution by institution basis.

2 Background and Institutional Setting

2.1 The recruitment sequence

Structured recruitment proceeds through a fixed sequence of rounds. In each round, PNMs visit a subset of chapters, and both sides submit ranked lists. A matching procedure assigns each PNM to the chapters that will invite her back, subject to the release figures. The number of visits per PNM contracts across rounds, commonly from all chapters in the opening round to two in the preference round. The final round produces a bid list, and PNMs who accept are placed.

Three centrally determined quantities govern the process. **Quota** is the number of new members each chapter may pledge, computed from the number of PNMs remaining at the preference round divided by the number of participating chapters; it is uniform across chapters on a given campus. **Total** is the maximum permitted chapter size, conventionally set at or near the size of the largest chapter, which permits smaller chapters to recruit toward parity. **Release figures** are the per-chapter, per-round invitation counts produced by RFM.

2.2 Why release figures are not uniform

A chapter cannot fill quota from its invitation list alone, because PNMs it invites may prefer another chapter. RFM therefore inflates release figures above quota by an amount reflecting anticipated attrition. A chapter whose invitations are usually reciprocated needs a smaller inflation factor than one whose invitations frequently are not. This is a sound response to a real problem. Its distributional consequence is the subject of Section 4.

3 Formal Model

3.1 Primitives

Let $C = \{1, \dots, m\}$ index chapters and $P = \{1, \dots, n\}$ index PNMs entering recruitment. Let $t \in \{1, \dots, T\}$ index rounds, with T the preference round. Let v_t denote the number of chapters each PNM visits in round t, with $v_1 \geq v_2 \geq \cdots \geq v_T$.

Each PNM p holds a preference relation $\succ_p$ over C; each chapter c holds a preference relation $\succ_c$ over P. Let $r = (r_1, \dots, r_m) \in \mathbb{Z}^m_{\geq 0}$ denote the release figure vector for a given round.

Definition 1 (Round matching)**.** Given release figures r and submitted preferences, a round matching is a correspondence $\mu : C \to 2^P$ satisfying $|\mu(c)| \leq r_c$ for all c and $|\{c : p \in \mu(c)\}| \leq v_{t+1}$ for all p.

Definition 2 (Placement)**.** A PNM is *placed* if she receives and accepts a bid at the conclusion of round T. Let $N(r)$ denote the realized number of placements, and $n_c(r)$ the number of new members pledged by chapter c.

3.2 The current objective

Let q denote quota, τ denote total, and s_c the current size of chapter c. RFM as currently practiced solves

$$\begin{aligned} \max_{r \in \mathbb{Z}^m_{\geq 0}} \quad & \mathbb{E}[N(r)] \\ \text{s.t.} \quad & \mathbb{E}[n_c(r)] \leq q, \qquad \forall c \in C, \\ & s_c + \mathbb{E}[n_c(r)] \leq \tau, \quad \forall c \in C, \end{aligned} \tag{1}$$

with the expectation taken over the realized preference profiles, estimated in practice from historical reciprocity rates.

Remark 1 (The degeneracy that matters)**.** Problem (1) is typically *massively degenerate*. For realistic campus parameters, many distinct release figure vectors attain the same value of $\mathbb{E}[N(r)]$, or values within one or two placements of the optimum. These vectors are not equivalent from the chapters' perspective: they differ substantially in the induced distribution $(n_1, \dots, n_m)$. The current mechanism selects among them without a stated criterion. This degeneracy is the opening that Section 5 exploits.

3.3 Distributional statistics

Define chapter fill rate $\phi_c(r) = n_c(r)/q$ and let $\bar{\phi}(r)$ be its mean across chapters. We use two dispersion measures, the coefficient of variation and the Gini coefficient of fill rate:

$$\mathrm{CV}(r) = \frac{1}{\bar{\phi}(r)} \sqrt{\frac{1}{m} \sum_{c \in C} \big(\phi_c(r) - \bar{\phi}(r)\big)^2}, \tag{2}$$

$$G(r) = \frac{\sum_{c \in C} \sum_{c' \in C} \big|\phi_c(r) - \phi_{c'}(r)\big|}{2m^2\, \bar{\phi}(r)}. \tag{3}$$

Neither appears in (1).

4 Four Structural Failure Modes

4.1 F1: Demand-derived release figures concentrate attrition risk

Because release figures scale inversely with reciprocity, a low-demand chapter receives a larger release figure and therefore invites back more PNMs per available slot. The composition of that larger invitation pool is the problem: it is disproportionately populated by PNMs who were released by their higher-preference chapters. The chapter absorbs greater attrition risk per invitation and per slot, and its recruitment workforce is spread across more conversations of lower conversion probability. The correction for low demand is, at the margin, also a tax on it.

4.2 F2: Uniform quota against non-uniform demand

Quota is a single number applied to every chapter. A chapter operating at half the demand of the campus leader is nonetheless evaluated against the same target, and is designated as having failed quota when it does not reach it. The designation is not merely descriptive. It is observable to PNMs, to alumnae, and to inter/national headquarters, and it enters the following year's demand formation.

4.3 F3: The objective is single-term and terminates at Bid Day

Two consequences follow. First, as established in Remark 1, dispersion is unconstrained, so the mechanism is free to select highly unequal allocations that are optimal or near-optimal in aggregate. Second, the objective counts bids accepted, not members retained. A bid accepted and depledged within six weeks scores identically to a bid that reaches initiation. Given that the marginal PNM placed into a chapter she ranked last is precisely the PNM most likely to

depledge, the mechanism can report success on placements it does not deliver in membership.

4.4 F4: Year-over-year compounding

The failure modes above are not static. Let $d_c^{(t)}$ denote demand for chapter c in year t. The empirical regularity is a positive feedback loop: demand $d_c^{(t)}$ determines the pledge class $n_c^{(t)}$, which determines chapter size $s_c^{(t+1)}$, which determines available recruiters, dues revenue, and campus visibility, which in turn shape $d_c^{(t+1)}$:

$$d_c^{(t)} \longrightarrow n_c^{(t)} \longrightarrow s_c^{(t+1)} \longrightarrow \text{capacity} \longrightarrow d_c^{(t+1)}. \quad (4)$$

Because RFM takes $d_c^{(t)}$ as an input and returns $n_c^{(t)}$ as an output, it sits inside this loop as an amplifier. A single poor recruitment year propagates into the next year's inputs without damping. The mechanism has no term that resists this propagation.

5 Proposed Reformulation

We propose six changes, ordered by expected benefit per unit of implementation cost. R1 and R4 are the highest-leverage and require no change to the recruitment experience of PNMs. Table 1 maps each change to the failure modes it targets.

5.1 R1: ε-constrained lexicographic objective

Rather than adding a fairness penalty to the objective, which requires committing to an exchange rate between placements and equity, we exploit the degeneracy identified in Remark 1. Let

$$N^\star = \max_r \mathbb{E}[N(r)], \qquad \mathcal{R}_\varepsilon = \left\{r : \mathbb{E}[N(r)] \geq (1-\varepsilon)N^\star\right\}. \quad (5)$$

The reformulated mechanism solves

$$\begin{aligned} \min_{r \in \mathcal{R}_\varepsilon} \quad & \mathrm{CV}(r) \\ \text{s.t.} \quad & \underline{q} \leq \mathbb{E}[n_c(r)] \leq \overline{q}, \quad \forall c \in C, \\ & s_c + \mathbb{E}[n_c(r)] \leq \tau, \quad \forall c \in C, \end{aligned} \quad (6)$$

with ε a small tolerance set by the Panhellenic council, on the order of 0.01 to 0.02. The interpretation is direct: among all allocations that place at least $(1-\varepsilon)$ of the maximum number of PNMs, select the one that distributes fill rate most evenly. Where the degeneracy is exact, $\varepsilon = 0$ already yields equity gains at zero placement cost.

An equivalent formulation replaces the lexicographic structure with an inequality-averse social welfare function in the spirit of Atkinson [2]:

$$W_\alpha(r) = \begin{cases} \dfrac{1}{1-\alpha} \displaystyle\sum_{c \in C} \phi_c(r)^{1-\alpha}, & \alpha \neq 1, \\ \displaystyle\sum_{c \in C} \log \phi_c(r), & \alpha = 1, \end{cases} \quad (7)$$

where $\alpha = 0$ recovers the current objective and $\alpha \to \infty$ approaches maximin. Reporting outcomes across a range of α makes the council's equity preference an explicit, auditable parameter rather than an artifact of solver tie-breaking. The tradeoff this exposes is the price of fairness in the sense of Bertsimas et al. [3].

5.2 R2: Retention weighting

Replace the placement count with a retention-weighted count. Let $\pi_{cp} \in [0,1]$ denote the estimated probability that PNM p, if placed in chapter c, is retained through initiation. Then

$$\tilde{N}(r) = \sum_{c \in C} \sum_{p \in \mu(c)} \pi_{cp}, \qquad \tilde{n}_c(r) = \sum_{p \in \mu(c)} \pi_{cp}. \quad (8)$$

The quantity π_{cp} is estimable from historical cohorts using features available at allocation time: the PNM's stated rank of c, round of first mutual match, whether c was a snap bid or continuous open bidding placement, and chapter-level historical retention. A regularized logistic model is adequate and preferable to a higher-capacity model here, because the coefficients must be interpretable to a governing council.

The effect on struggling chapters is direct. The current objective rewards configurations that fill a chapter with PNMs who will leave. The retention-weighted objective does not, and therefore stops treating a fragile class as equivalent to a durable one.

5.3 R3: Cardinal preference elicitation

Ordinal preference cards cannot distinguish a PNM who ranks a chapter second with genuine enthusiasm from one who ranks it second because the card required two entries. This is an information loss that falls asymmetrically on chapters outside the top tier, whose genuine supporters are rendered indistinguishable from indifferent ones.

We propose a constrained point allocation: each PNM distributes a fixed budget B across the chapters she visited, subject to $\sum_c b_{pc} = B$ and $0 \leq b_{pc} \leq b_{\max}$. The cap $b_{\max}$ limits the strategic value of concentrating the full budget on a single chapter. The resulting cardinal utilities improve the estimate of $\mathbb{E}[N(r)]$ and of π_{cp}, and give low-demand chapters credit for intensity that ordinal data erases.

Strategic caveat. Cardinal elicitation is not strategyproof, and the literature on tie-breaking and preference reporting in school choice [1, 4] establishes that reporting incentives interact with mechanism structure in ways that are not always intuitive. R3 should therefore be evaluated under strategic as well as truthful reporting in simulation (Section 6) before adoption, and is the change we recommend deploying last.

5.4 R4: Temporal smoothing of release figures

To damp the loop in (4), constrain the year-over-year change in each chapter's effective release figure:

$$r_c^{(t)} = \mathrm{clip}\Big(\hat{r}_c^{(t)},\ \underline{\rho}\, r_c^{(t-1)},\ \overline{\rho}\, r_c^{(t-1)}\Big), \quad (9)$$

where $\hat{r}_c^{(t)}$ is the unsmoothed figure computed from current-year demand and $\underline{\rho} < 1 < \overline{\rho}$ are policy-set bounds, for example $\underline{\rho} = 0.85$ and $\overline{\rho} = 1.15$. Equivalently, demand inputs may be smoothed by exponential weighting, $\bar{d}_c^{(t)} = \lambda d_c^{(t)} + (1-\lambda)\bar{d}_c^{(t-1)}$.

The design intent is that a single anomalous year, which may reflect a leadership transition, a facility closure, or a transient reputational event, does not fully propagate into the following year's allocation. R4 is the only proposed change that acts directly on

the compounding mechanism rather than on a single year's cross-section.

5.5 R5: Quota as a band

Replace point quota q with a band $[\underline{q}, \overline{q}]$, where $\underline{q} = \lfloor (1-\delta)q \rfloor$ and $\overline{q} = \lceil (1+\delta)q \rceil$. A chapter in a declared rebuilding status may petition in advance for a target at the lower end of the band without incurring the failed-quota designation.

Quota additions already accomplish something similar, but reactively and after the fact. Making the adjustment prospective and procedural removes the public failure signal that currently enters the following year's demand formation through (4), without changing the arithmetic of who is placed where.

5.6 R6: Median-anchored total with a dispersion trigger

Total anchored to the largest chapter permits the highest-demand chapters to absorb the most members and correspondingly shrinks the pool available to others. We propose anchoring at or near the median chapter size,

$$\tau^{(t)} = \text{median}_{c \in C} \left(s_c^{(t)} \right) + \kappa, \tag{10}$$

with a small headroom constant κ, together with an automatic trigger: if $\text{CV}(s^{(t)})$ exceeds a threshold θ, then τ adjusts downward by a preset increment for the following cycle. This makes total a stabilizing control variable rather than a passive record of the largest chapter's success. R6 has the largest expected effect on dispersion and the largest expected political cost, since it binds most on the chapters with the most institutional influence.

5.7 Complementary operational change

Continuous open bidding (COB) currently functions as a relief valve and is treated administratively as remedial, which is a substantial part of why it underperforms. We recommend funding COB, scheduling it on the Panhellenic calendar, marketing it council-wide, and reporting its results as ordinary recruitment rather than as a deficiency measure. This is not a change to the formula, but it changes the distribution the formula is trying to correct.

6 Evaluation Protocol

No change proposed here should be deployed on a live recruitment before retrospective and simulated evaluation. Recruitment is annual and non-repeatable; an error is not correctable within the cycle in which it occurs.

6.1 Phase 1: Historical backtesting

For each participating campus, assemble three to five years of round-level records: submitted preference lists, release figures, round matchings, bid acceptances, and retention through initiation. Recompute release figures under each proposed formulation and compare the counterfactual allocation against the realized one. Backtesting bounds the plausible effect size but cannot capture behavioral response, since PNM preferences are held fixed at their historical values.

Algorithm 1 Policy evaluation sweep

Input: policies $\mathcal{P}$, seeds Ω, model $\mathcal{M}$
for each policy $\Pi \in \mathcal{P}$ **do**
  for each seed $\omega \in \Omega$ **do**
    Draw profiles $(\succ_P, \succ_C) \sim \mathcal{M}(\omega)$
    for $t = 1$ to T **do**
      $r^{(t)} \leftarrow \Pi\big(\hat{d},\, q,\, \tau,\, r^{(t-1)}\big)$
      $\mu^{(t)} \leftarrow \textsc{RoundMatch}\big(r^{(t)}, \succ_P, \succ_C\big)$
    end for
    Record $\big(N,\ \tilde{N},\ \text{CV},\ G,\ \min_c \phi_c\big)$
  end for
end for
Return outcome distribution per policy

6.2 Phase 2: Agent-based counterfactual simulation

Construct a simulator with PNM agents whose preferences are drawn from distributions calibrated to the historical data, including a common-value component reflecting shared campus reputation and an idiosyncratic component. For each policy and each seeded scenario, simulate the full recruitment sequence and record the outcome vector, as in Algorithm 1. Vary the strength of the common-value component to test sensitivity to demand concentration, and simulate both truthful and strategic reporting under R3.

6.3 Metrics

We report aggregate placement N, the count the current mechanism optimizes; retention-weighted placement $\tilde{N}$, the placements surviving to initiation; the dispersion measures CV and G over chapter fill rate; the floor $\min_c \phi_c$, which is the worst-off chapter's fill rate; the price of fairness $\text{PoF} = (N^\star - N^\Pi)/N^\star$, the aggregate placement forgone under policy Π relative to the unconstrained optimum [3]; and stability, measured as the rate of blocking pairs relative to the baseline mechanism, following standard definitions in two-sided matching [7].

6.4 Hypotheses

H1
Under R1 with $\varepsilon = 0.01$, CV decreases materially relative to baseline while $\text{PoF} \le 0.01$ by construction. *Falsified if* the feasible set $\mathcal{R}_\varepsilon$ is near-singleton in realistic campus configurations, that is, if the degeneracy of Remark 1 does not obtain.

H2
Under R2, $\tilde{N}$ increases relative to baseline even where N decreases slightly. *Falsified if* π_{cp} is not predictable above base rate from allocation-time features.

H3
Under R4, the variance of ϕ_c across consecutive years decreases, and the autocorrelation of chapter-level demand shocks weakens. *Falsified if* demand shocks are predominantly persistent rather than transient, in which case smoothing merely delays adjustment.

H4
Under R6, $\min_c \phi_c$ increases and CV decreases, with PoF bounded by the fraction of PNMs whose acceptable set is confined to the

Table 1: Proposed changes mapped to the failure modes of Section 4.

Change	Targets	Implementation cost	Expected dispersion effect
R1 ε-constrained objective	F3	Low (solver only)	Moderate
R2 Retention weighting	F3	Moderate (model and data)	Moderate
R3 Cardinal elicitation	F1	High (PNM-facing change)	Low to moderate
R4 Temporal smoothing	F4	Low (input processing)	Moderate to high
R5 Quota banding	F2	Low (policy only)	Indirect
R6 Median-anchored total	F1, F2	Moderate (policy, contested)	High

largest chapters. *Falsified if* PNMs released from top chapters withdraw rather than accept alternative bids at a rate that offsets the redistribution.

H4 identifies the central empirical uncertainty of this paper. Redistribution through total only improves outcomes if the redistributed PNMs remain in the process. If they instead withdraw, R6 reduces aggregate placement without improving anyone's fill rate. This parameter, the withdrawal elasticity with respect to release from a top-preference chapter, should be estimated before R6 is adopted anywhere.

7 Limitations and Deployment Ethics

Demand is exogenous. None of R1 through R6 makes a chapter more desirable. They redistribute the consequences of existing preferences and damp the mechanism's amplification of them. A chapter with a genuine and persistent programming or reputational deficit will remain a low-demand chapter under any allocation rule.

Fairness is a contested choice, not a computed one. The parameters ε, α, δ, and θ encode a normative position on how much aggregate placement is worth trading for distributional equality. This paper does not argue for particular values. It argues that the values should be chosen deliberately by a governing body and published, rather than emerging implicitly from solver behavior.

Data governance. The retention model of R2 uses individual-level student records. Any implementation must satisfy institutional review, comply with applicable student-records regulation, restrict features to those defensible on face validity, and audit for disparate impact across protected characteristics. A model that improves aggregate retention while systematically down-weighting an identifiable subgroup of PNMs is not an acceptable outcome, and it will not be detected by the metrics of Section 6 unless subgroup analysis is specified in advance.

Generality. Campuses differ in chapter count, demand concentration, housing constraints, and institutional support. The evaluation protocol should be run per campus. We expect the effect sizes of R1 through R6 to depend strongly on baseline demand concentration, with the largest gains where dispersion is already high.

8 Conclusion

Release Figure Methodology solves the problem it was designed to solve. Aggregate placement is high, and the mechanism is a clear improvement over the practices it replaced. Its limitation is not an error but an omission: the objective function contains one term, and the distribution of outcomes across chapters is therefore determined as a residual rather than as a decision.

We have argued that this omission has consequences, that those consequences compound across years, and that a substantial portion of the resulting inequality can be reduced at little or no cost to aggregate placement, because the underlying optimization problem is degenerate and the mechanism currently breaks ties without a criterion. The six changes proposed here make the equity and efficiency tradeoff explicit, parameterize it, and place the choice of parameters where it belongs, with the governing council rather than with the solver.

The immediate next step is empirical. Phase 1 backtesting requires no policy change, no consent from chapters, and no risk to any recruitment cycle. It will either confirm that $\mathcal{R}_\varepsilon$ contains meaningfully more equitable allocations at negligible placement cost, or it will falsify H1 and substantially weaken the argument of this paper. Either outcome is worth knowing before the next recruitment.

References

[1] Atila Abdulkadiroğlu, Parag A. Pathak, and Alvin E. Roth. 2009. Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match. *American Economic Review* 99, 5 (2009), 1954–1978.

[2] Anthony B. Atkinson. 1970. On the Measurement of Inequality. *Journal of Economic Theory* 2, 3 (1970), 244–263.

[3] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. 2011. The Price of Fairness. *Operations Research* 59, 1 (2011), 17–31.

[4] Aytek Erdil and Haluk Ergin. 2008. What's the Matter with Tie-Breaking? Improving Efficiency in School Choice. *American Economic Review* 98, 3 (2008), 669–689.

[5] David Gale and Lloyd S. Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69, 1 (1962), 9–15.

[6] Alvin E. Roth. 2008. What Have We Learned from Market Design? *The Economic Journal* 118, 527 (2008), 285–310.

[7] Alvin E. Roth and Marilda Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis.* Cambridge University Press, Cambridge, UK.