

Application of Market Design in FRC Drive Team Selection: Optimization Using the Gale-Shapley Algorithm

Yao Chen

RDFZ Xishan School, Beijing, China
chenyao080923@gmail.com

Abstract. In the FIRST Robotics Competition (FRC), driver performance is a critical determinant of a team's success. However, the traditional selection process for drive teams is often subjective and inefficient, relying heavily on the personal impressions of mentors or simplistic skill rankings. This study applies market design, specifically the Gale-Shapley algorithm, to optimize the driver selection process. By utilizing evaluation criteria established by elite teams and official FIRST guidelines—including communication, technical talent, and composure—this study builds a quantitative scoring system. These scores are converted into structured "preference lists" to simulate a two-sided matching process, mirroring the heuristic approach of the CATME Team-Maker model. Through Python-based algorithmic simulation involving 20 students competing for multiple drive team roles, the results demonstrate that the Gale-Shapley algorithm effectively aligns student preferences with the operational needs of the team. Compared to traditional manual selection, the algorithm doubled top-choice satisfaction and eliminated unstable blocking pairs. Ultimately, this approach replaces subjective biases with a transparent and highly efficient selection mechanism. This study creatively applies the established mathematical algorithm (Gale-Shapley) to the specific domain of FRC, achieving a practical transformation from theory to application.

Keywords: Market Design, Gale-Shapley Algorithm, FRC Driver Selection, Stable Matching, Team Formation

1. Introduction

The FRC represents a pinnacle of high school STEM engagement, challenging students on both engineering ingenuity and strategic execution. Within this multifaceted contest, the composition and performance of the drive team, which is the core unit responsible for operating the robot during matches, are widely recognized as a decisive factor in competitive success [1, 2].

Despite its importance, the current methods for selecting drive team members within many high school FRC teams are often flawed. Traditionally, selection relies either on the subjective impressions of mentors or on simple, one-dimensional driving test rankings. This approach overlooks a crucial dynamic: the specific compatibility between a candidate's distinct traits and the unique requirements of different roles. Furthermore, traditional methods rarely account for the personal preferences and psychological readiness of the students themselves, leading to suboptimal

team synergy and decreased passion. This presents a significant research gap. While engineering and strategy in FRC have been heavily optimized through data, human resource management lacks an algorithmic approach. To address this issue, this study introduces the two-sided matching theory from market design into the context of FRC [3, 4].

Methodologically, this study employs a literature analysis to establish robust evaluation metrics and utilizes mathematical modeling to construct a student-role matching model based on the Gale-Shapley (Deferred Acceptance) algorithm [5]. This research provides FRC teams with a scientific decision-making tool to enhance team efficiency and expand the application boundaries of matching algorithms to small-scale, skill-based team formations.

2. FRC driver selection criteria and data foundation

To construct an effective algorithmic matching model, it is first necessary to translate abstract human qualities into measurable data. The drive team in FRC typically consists of four specialized roles: the Driver, the Operator, the Human Player, and the Technician. Elite FRC programs, such as Team 254 (The Cheesy Poofs), have established comprehensive criteria emphasizing that soft skills are just as vital as technical manipulation [6]. Based on their handbook and official FIRST guidelines [7], this study identifies and defines three core competencies required for a successful drive team member:

(1) Communication: The ability to effectively and respectfully convey new ideas and strategies under the high-stress environment of a match. This competency is selected because the official FIRST guide [7] and Team 254's handbook [6] both rank it as a top-priority skill. Communication is especially critical for the Drive Coach and Technician roles, whose primary functions during a match involve relaying strategic calls and coordinating repairs, respectively. However, all four roles require baseline communication to maintain team cohesion.

(2) Technical Talent vs. Experience: Hand-eye coordination, spatial awareness, and a deep understanding of the robot's mechanical limits. This dimension is drawn directly from the FIRST Chesapeake guide's discussion of "Talent vs. Experience" [7], which argues that a candidate's raw, innate ability to operate the robot is "a powerful differentiator." This competency carries the highest weight for the Driver and Operator roles, whose primary task is the physical manipulation of the robot, while it is relatively less critical for the Drive Coach.

(3) Composure: Emotional stability and the capacity to coolly evaluate the situation when facing unexpected robot malfunctions or intense competitive pressure. The official FIRST documentation [1] highlights that FRC tournaments are complex environments with immense time pressure. This competency is prioritized for the Driver and Drive Coach, who must make split-second decisions under duress, but it also forms a necessary baseline for the Human Player, who performs under direct audience observation.

In this model, mentors assess candidates across these three dimensions using a standardized rubric (e.g., a 1-10 scale for each competency). These distinct sub-scores are then aggregated using weighted averages tailored to each specific role. For instance, the "Driver" Credit Score may weight Technical Talent at 50%, Composure at 30%, and Communication at 20%, reflecting the Driver's need for superior control skills, whereas the "Drive Coach" Credit Score may weight Communication at 50%, Composure at 30%, and Technical Talent at 20%. The resulting composite score for each candidate-role pair is referred to as a "Credit Score"—a single numerical value that quantifies a candidate's overall suitability for a particular role. This structured data foundation is essential for generating the strict preference lists required by stable matching models [8]. It is important to note that the specific weights used in this study are illustrative and were informed by

the relative emphasis placed on each skill in the Team 254 handbook [6] and official FIRST guidelines [1, 7]; in practice, individual teams should calibrate these weights to their own strategic priorities.

3. Theoretical framework: market design and algorithm adaptation

3.1. Two-sided matching mechanisms

The fundamental theory behind our model is the concept of two-sided matching, a cornerstone of Market Design. Originally introduced by Gale and Shapley in 1962, the stable matching theory was designed to solve the college admissions problem [5]. The core objective of the Gale-Shapley Deferred Acceptance (DA) algorithm is to pair agents from two disjoint sets based on their ranked preferences, ensuring a "stable" outcome.

A match is defined as "stable" if there is no blocking pair—meaning there are no two entities (a student and a role) who would both prefer each other over their current assignments [3]. The deferred acceptance mechanism has since been widely implemented in various large-scale allocation systems, most notably the National Residency Matching Program (NRMP) for medical students in the United States. In the context of FRC, the two sides of the market are the "Student Candidates" and the "Drive Team Roles."

3.2. Insights from the CATME model

While the Gale-Shapley algorithm described in Section 3.1 was originally designed for large-scale markets such as college admissions and medical residencies, its adaptation to small-group, skill-based team formation requires a practical bridge. The CATME Team-Maker model provides precisely this bridge. Meulbroek et al. [9] demonstrated how to integrate the Gale-Shapley algorithm into an existing team formation tool to maximize student passion and team cohesion. In their capstone course application, they treated student project preferences as one side of the market and instructor or sponsor evaluations as the other, showing that over 90% of students could be placed within their top three choices [9].

Inspired by this approach, our FRC selection model directly maps the two sides of the matching market: the "student side" is represented by each candidate's ranked preference list for the four roles, while the "role side" is represented by the preference list generated from the mentor-evaluated Credit Scores, which objectively reflects each role's skill requirements. By mathematically balancing soft skills and personal preferences, the algorithm becomes perfectly applicable to the high-stakes environment of FRC driver selection.

4. Model construction and algorithm implementation

4.1. Generating preference lists

To run the matching algorithm, the system requires two complete sets of preference lists.

Student Preference Lists: Every student candidate submits a ranked list of the drive team roles they wish to fulfill, ordered from their most preferred (Rank 1) to least preferred (Rank 4).

Role Preference Lists (Team Needs): Unlike human students, the "Roles" cannot express subjective preference. Therefore, the preference list for each role is generated automatically based on the candidates' Credit Scores. For example, the "Driver" role will rank all candidates strictly in descending order of their Driver Credit Score.

Since all slots for a given role (e.g., all 5 Driver slots) are considered identical in terms of requirements, they share a single preference list ranked by the role-specific Credit Score.

4.2. Algorithm process

With the preference lists established, the study utilizes a modified student-proposing Deferred Acceptance algorithm. The step-by-step process is as follows:

1. Initialization: All students and roles begin as "unmatched."
2. Proposal Phase: Each unmatched student proposes to their top-choice role.
3. Tentative Acceptance Phase: For each role, all current proposals are compared. The role tentatively accepts the top-ranked students (according to its preference list) up to its capacity limit (e.g., 5 students). Any proposing student not among the top-ranked is rejected.
4. Iteration: The rejected students then propose to their next highest preferred role. The roles evaluate the new proposals against their provisionally accepted students. If a new proposal comes from a better-qualified candidate, the role rejects the previously held student and provisionally accepts the new candidate.
5. Termination: This process repeats until no student is left unmatched.

5. Case study and discussion

To test the viability of this algorithmic solution, a Python-based simulation was conducted mirroring a realistic FRC team size. Often, large organizations like Team 254 manage multiple FTC sub-teams simultaneously (e.g., Teams 23222, 23223, 23224) [6], requiring the assignment of many students to identical roles across different teams. The simulation involved 20 student candidates competing for 4 drive team roles, with each role having a capacity of 5 slots (representing 5 complete drive teams).

The students were assigned randomized but realistic Credit Scores across the four roles, and each student generated a complete preference ranking. The Gale-Shapley (GS) algorithm was executed alongside a simulated "Manual Selection" method, which mimicked traditional mentor-led selection by assigning students sequentially based purely on their highest single skill score. Utilizing metrics proposed by Serek and Zhaparov [10], the study evaluated "preference satisfaction" and "workload/score efficiency."

As shown in Table 1, the GS algorithm placed 14 out of 20 students into their 1st-choice role, compared to only 7 under Manual Selection.

Table 1. Comparison of student preference satisfaction between gale-shapley and manual selection methods

Metric	1st Choice	2nd Choice	3rd Choice	4th Choice
Gale-Shapley (GS)	14	5	1	0
Manual Selection	7	3	6	4

As shown in Table 2, the Manual Selection method achieved a marginally higher average Credit Score but at the cost of significant matching instability.

Table 2. Comparison of team credit scores and matching stability

Assignment Model	Average Team Credit Score	Unstable Matches (Blocking Pairs)
Gale-Shapley (GS)	82.3	0
Manual Selection	90.3	> 5

The contrast between the two methods is further visualized in Figure 1. As shown in Figure 1, the GS algorithm demonstrates a strong concentration of placements at the 1st-choice level (14 students), while Manual Selection yields a more dispersed and less favorable distribution, with a substantial number of students pushed into their 3rd and 4th choices.




Figure 1. Distribution of student role assignments by preference rank under gale-shapley vs. manual selection

The simulation results revealed a fascinating trade-off:

1. Preference Satisfaction: Under the GS algorithm, 14 out of 20 students (70%) received their 1st choice role, and 95% received their top two choices. In contrast, the Manual Selection method only granted 7 students (35%) their 1st choice, forcing 20% of the students into their least desired 4th choice role.

2. Stability and Fairness: The GS output contained zero blocking pairs, guaranteeing procedural transparency and ensuring that no student felt unjustly denied a role they preferred. This result directly validates the core theoretical guarantee of the Gale-Shapley algorithm as described in Section 3.1: by achieving a stable matching with zero blocking pairs, the algorithm successfully fulfilled its fundamental design objective of eliminating the structural instability that could otherwise lead to internal dissatisfaction, perceived unfairness, and ultimately personnel turnover within the drive team. In contrast, the Manual Selection method produced more than five blocking pairs, meaning there existed multiple student-role combinations where both the student and the role would have mutually preferred each other over their current assignment—a condition that is a recognized source of grievance in team settings [3, 5].

3. Score Efficiency: The Manual Selection method achieved a higher Average Team Credit Score (90.3 vs. 82.3), as shown in Table 2. However, this came at the severe cost of student dissatisfaction. The GS algorithm naturally cascades talent to optimize the overall baseline capability while dramatically improving student morale.

6. Conclusion

This study demonstrates that the Gale-Shapley algorithm, when combined with a quantitative evaluation criteria framework, can generate an optimal, stable, and transparent FRC drive team roster. By utilizing the Deferred Acceptance mechanism, the proposed model effectively avoids the inherent biases and emotional friction associated with traditional subjective selection processes. It

ensures that both student aspirations and team performance requirements are symmetrically optimized.

However, certain limitations must be acknowledged. First, algorithmic models struggle to fully capture the complex psychological dynamics and interpersonal "chemistry" between specific team members, which are often qualitative and difficult to score. Furthermore, the model's effectiveness depends almost entirely on the accuracy and objectivity of the initial input Credit Scores provided by mentors. In real-world deployment, several practical challenges may arise. Collecting truthful and reliable preference rankings from students presents its own difficulties: students may engage in strategic behavior—such as not ranking their true first choice if they perceive fierce competition for that role—despite the theoretical strategy-proofness of the Deferred Acceptance mechanism for the proposing side.

Future research could introduce machine learning algorithms to assist in the initial scoring phase, minimizing human bias in skill evaluation. Similarly, computer vision systems could be used to evaluate students' reaction speed, hand-eye coordination accuracy, and composure during simulated robot operation drills, providing quantitative inputs for both the Technical Talent and Composure dimensions.

References

- [1] FIRST Robotics Competition. (2025). Selecting drive team members. FIRST Official Resource Library. <https://www.firstinspires.org/resource-library/frc/selecting-drive-team-members>
- [2] Fernbank LINKS, & Columbus Space Program. (2017). The FRC survival guide. Chief Delphi. <https://www.chiefdelphi.com/t/the-frc-survival-guide/158976>
- [3] Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. *The American Economic Review*, 93(3), 729–747. <https://doi.org/10.1257/00028280322157061>
- [4] Pathak, P. A. (2011). The mechanism design approach to student assignment. *Annual Review of Economics*, 3(1), 513–536. <https://doi.org/10.1146/annurev-economics-061109-080213>
- [5] Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1), 9–15. <https://doi.org/10.1080/00029890.1962.11989827>
- [6] Team 254. (2024). Bellarmine College Preparatory Robotics Student Handbook. Team 254.
- [7] FIRST Chesapeake. (2024, October 25). Selecting drivers for FRC. FIRST Inspires. <https://www.firstchesapeake.org/post/selecting-drivers-for-frc>
- [8] Perach, N., & Anily, S. (2022). Stable matching of student-groups to dormitories. *European Journal of Operational Research*, 302(1), 50–61. <https://doi.org/10.1016/j.ejor.2021.12.048>
- [9] Meulbroek, D., Ferguson, D., Ohland, M., & Berry, F. (2019). Forming more effective teams using CATME Team-Maker and the Gale-Shapley algorithm. In 2019 IEEE Global Engineering Education Conference (EDUCON) (pp. 1196–1200). IEEE. <https://doi.org/10.1109/EDUCON.2019.8725173>
- [10] Serek, A., & Zhaparov, M. (2023). Algorithm comparison for student-supervisor matching in supervisorship system development: K-Means vs. one-to-many Gale-Shapley. *Information Sciences Letters*, 12(12), 2417–2425. <https://doi.org/10.18576/isl/121205>