Personal Statement

Peng Shi
Assistant Professor, Operations Management Group
USC Marshall School of Business — May 2025

My research investigates a fundamental question with far-reaching societal implications: **How should we design systems that match individuals based on their preferences?** This question drives critical decisions in domains ranging from public school assignment and affordable housing allocation to the algorithms powering digital marketplaces that connect millions of consumers with service providers daily. In each setting, preferences are diverse, stakes are high, and design choices directly shape opportunity and welfare for countless individuals.

Despite the importance of these systems, practitioners often operate with limited visibility into critical metrics. Preferences are difficult to observe directly, and the effects of alternative policies are challenging to estimate without disrupting existing operations. Consequently, system designers frequently rely on intuition, local experimentation, and heuristic adjustments—approaches that may overlook superior solutions or generate unintended consequences.

This is precisely where my research makes significant contributions: I develop mathematical models that rigorously characterize the design space, identify provably near-optimal solutions, and provide actionable guidance even with limited data. My work bridges theory and practice through two complementary research streams:

1. **Matchmaking Systems for Scarce Public Resources:** My research on school choice, subsidized housing, and organ allocation demonstrates how optimization can enhance access to life-changing opportunities while balancing competing policy objectives.
2. **Two-Sided Digital Marketplaces:** I analyze how platforms like Amazon, Angi, Airbnb, Google, Yelp, and others can efficiently facilitate billions of matches between consumers and providers, where seemingly small design changes can have system-wide effects.

In both domains, I create models that are analytically rigorous yet practical—frameworks that help stakeholders understand their systems, evaluate policy trade-offs, and design better mechanisms. My contributions have shaped policy reform in Boston's public school system, influenced platform policies in digital marketplaces, and earned recognition through awards such as the 2024 Frederick W. Lanchester Prize. The sections that follow detail these contributions and outline my vision for future research.

1. Matchmaking Systems for Scarce Public Resources

A central theme of my research has been improving the design of public-sector matching systems where individuals are assigned to limited resources based on their preferences. While the economics literature has deeply analyzed school choice mechanisms, prior work has almost universally treated key policy levers—such as choice menus and priority rules—as fixed parameters. In contrast, my research demonstrates that these levers can be systematically optimized to better serve diverse stakeholders.

This research began with a collaboration with Boston Public Schools, which in 2012-2013 sought to replace a decades-old assignment system based on three large geographic zones. Policymakers needed a design that balanced equity of access, shorter commutes, and neighborhood cohesion. In "**Guiding**

School-Choice Reform through Novel Application of Operations Research" (*Interfaces*, 2015), I worked with the district to evaluate and enhance proposed assignment plans. I developed a multinomial logit (MNL) discrete choice model using historical ranking data to simulate how families would respond to each proposal and to assess resulting assignment outcomes. These simulations enabled city officials to compare metrics like access to quality schools, travel distance, and neighborhood cohesion across different plans.

The outcome of this collaboration—the "Home-Based Plan," which created personalized choice menus centered around each student's home—was adopted in 2014 and remains in use today. This reform received coverage in both the *New York Times*[1] and *Boston Globe*[2] and marked the first application of simulation-based optimization to school choice menu design.

To move beyond evaluating fixed proposals, I developed a tractable optimization framework to search the vast space of possible designs. In "**Optimal Allocation without Money: an Engineering Approach**" (*Management Science*, 2015), I established that under mild assumptions, any assignment mechanism satisfying standard fairness properties can be replicated by the deferred acceptance (DA) algorithm with appropriately chosen inputs—specifically, choice menus and priority structures. This result justified using DA as a design scaffold and shifted the challenge to optimizing its inputs. I introduced an efficient algorithm to solve this optimization problem under MNL preferences, incorporating constraints on school capacity and average travel distance.

I extended this methodology in "**Optimal Priority-Based Allocation Mechanisms**" (*Management Science*, 2021) to accommodate more general preferences and policy constraints. By drawing a novel connection to assortment optimization in revenue management, I incorporated geographically defined bus service areas—a critical concern for Boston's transportation planning. The resulting plan outperformed the city's existing policy on multiple dimensions: match quality, access equity, predictability, transportation costs, and neighborhood cohesion. These papers demonstrate that assignment policies can be rigorously optimized using tools from discrete choice modeling and operations research.

One social objective emphasized by Boston policymakers was community cohesion, defined as the probability that students from the same neighborhood attend the same school. In "**Improving Community Cohesion in School Choice via Correlated-Lottery Implementation**" (*Operations Research*, 2014), I showed that this metric can be improved by introducing correlation in tie-breaking lottery numbers, without affecting any individual's marginal assignment probabilities. Using Boston data, I demonstrated that lottery correlation improved neighborhood cohesion more effectively than any of the menu or priority changes under consideration.

All of these models assume that family preferences can be approximated by a stable discrete choice model estimated from historical data. In "**How Well Do Structural Demand Models Work? Counterfactual Predictions in School Choice**" (*Journal of Econometrics*, 2021), co-authored with Parag Pathak, I tested this assumption by comparing the model's pre-reform predictions to actual post-reform outcomes under the Home-Based Plan. After controlling for demographic shifts, we found that the pre-reform model performed comparably to a re-estimated post-reform model, validating the stability of preferences and supporting the use of such models for policy design.

Beyond school choice, I analyzed matchmaking systems for subsidized housing, where applicants face dynamic trade-offs between accepting a current offer or waiting for a better one. In "**Design of Lotteries and Wait-lists for Affordable Housing Allocation**" (*Management Science*, 2020, with Nick Arnosti), we

[1] See https://www.nytimes.com/2013/03/15/education/boston-schools-adopt-new-placement-plan-for-students.html.
[2] See https://www.bostonglobe.com/metro/2012/10/27/boston-panel-winnow-options-for-school-assignment-changes-new-mit-proposal-leaps-top/jy8yIsqjLM23w9ECHXpknM/story.html.

developed a flexible model to analyze various real-world mechanisms—including independent lotteries used in New York City and waitlists used in Providence and Amsterdam. A key insight was that these seemingly different mechanisms can yield equivalent equilibrium outcomes. For instance, we showed that NYC's lottery is equivalent to Providence's waitlist, while Amsterdam's waitlist is equivalent to a system based on virtual currency. We further demonstrated that modest rule changes—such as restricting the number of lotteries an applicant may enter—could significantly improve efficiency. The paper was praised by Nobel Laureate Paul Milgrom as an example of "how mathematical analysis in the hands of skilled researchers can lend deep new insights into important practical problems."[3]

Collectively, these papers introduce tractable optimization models for public-sector matching systems that are both theoretically grounded and practically implementable.

2. Two-Sided Marketplaces

A central question in the design of two-sided marketplaces is how to maximize match quality while minimizing search and communication costs. While prior research has demonstrated the benefits of specific interventions, few studies characterize what is optimally achievable or offer general principles to guide platform structure. My research addresses this gap by developing a rigorous theoretical framework that identifies provably near-optimal matchmaking strategies and evaluates mechanism performance under real-world constraints.

In **"Optimal Matchmaking Strategy in Two-Sided Marketplaces"** (*Management Science*, 2022), I analyze four commonly used matchmaking strategies: (1) customer-initiated search (e.g., Google Maps, Houzz); (2) provider-initiated search (e.g., Bark, Porch); (3) bilateral search (e.g., Upwork, Yelp); and (4) platform-mediated recommendations (e.g., HomeAdvisor, Thumbtack). I demonstrate that regardless of marketplace parameters, the simple heuristic of "let the picky side pick" reliably identifies the optimal strategy among these four: when one side's preferences are more idiosyncratic and difficult to describe, that side should be enabled to search over potential partners and initiate contact. Furthermore, I prove that the resultant policy achieves near-optimal performance among all possible matchmaking strategies according to an information-theoretic metric, which quantifies matchmaking efficiency based on the communication required to achieve favorable outcomes. This result reveals that the four prevalent strategies, when implemented in markets with appropriate preference characteristics, are already near-optimal, even compared to hypothetically more complex mechanisms.

This analysis builds on my earlier work, **"Clearing Matching Markets Efficiently: Informative Signals and Match Recommendations"** (*Management Science*, 2019, with Ashlagi, Braverman, and Kanoria), which pioneered the application of information theory for the design of matchmaking strategies. While that paper established the core theoretical framework, it was constrained to settings with hard-to-describe preferences on both sides and exogenous pricing. My 2022 paper significantly extends this approach through novel modeling innovations that accommodate endogenous pricing and a broader spectrum of market environments, thereby better explaining the diverse matchmaking strategies observed in practice. The 2019 paper was honored with the prestigious 2024 Frederick W. Lanchester Prize, awarded by INFORMS for "the best contribution to operations research and management sciences published in English over the past five years." [4] The 2022 paper was subsequently highlighted by Management Science's Editor-in-Chief as both "theoretically elegant and practically insightful."[5]

[3] See https://www.informs.org/Blogs/ManSci-Blogs/Management-Science-Review/Effective-Allocation-of-Affordable-Housing.
[4] See https://www.informs.org/Recognizing-Excellence/INFORMS-Prizes/Frederick-W.-Lanchester-Prize.
[5] See https://connect.informs.org/discussion/management-science-march-2023.

While these two papers identify what is achievable through disparate matchmaking architectures, my more recent work examines how platforms should fine-tune their existing policy levers to optimize marketplace outcomes. In **"Optimal Match Recommendations in Two-Sided Marketplaces with Endogenous Prices"** (*Management Science*, 2024), I investigate platforms that recommend a limited set of providers to each customer and analyze how such platforms should design their recommendation algorithms to maximize social welfare—the total value generated, comprising customer surplus, provider profits, and platform revenue. A key methodological advancement over previous research is my incorporation of price endogeneity, where providers strategically set their prices outside of platform's direct control. This critical feature characterizes many real-world marketplaces but has historically eluded tractable analysis. Remarkably, I demonstrate that accounting for this strategic behavior actually simplifies the platform's optimization problem: a straightforward policy of ranking providers by predicted conversion rate achieves first-best social welfare. This means that the platform cannot improve welfare outcomes even if it were able to directly control providers' pricing decisions.

When accurate predictions of conversion rates are unavailable, I study alternative mechanisms based on pricing. In "**The Welfare Effects of Selling Leads in a Two-Sided Marketplace**" (major revision at *Management Science*), I analyze platforms like Angi, HomeAdvisor, and Thumbtack, which sell leads to providers. I demonstrate that charging market-clearing lead fees—determined by balancing supply and demand—guarantees at least 79.15% of the optimal social welfare achievable with any lead-pricing policy, and at least 58.19% of the first-best social welfare. This approach provides a robust, information-light benchmark that aligns with industry practice.

In "**Welfare-Optimal Policies for Sponsored Advertising in a Two-Sided Marketplace**" (*EC 2025*, journal submission in preparation), I examine platforms like Amazon, Google, and Yelp that combine organic rankings with paid sponsored ads. I show that ranking ads by the product of bid and clickthrough rate—a widely adopted industry heuristic—maximizes the sum of provider profits and ad revenue. To enhance customer surplus, the platform should implement a discount on ad fees based on estimated provider quality, with the magnitude calibrated to the reliability of these quality estimates. This quality-adjusted pricing policy achieves the best possible guarantee relative to the first-best social welfare. Notably, eliminating sponsored ads entirely and relying solely on organic rankings results in strictly worse guarantees, demonstrating that properly priced sponsored ads always benefit the marketplace. Conceptually, this paper bridges my earlier findings: it interpolates between settings with accurate preference prediction (where conversion rate ranking yields first-best outcomes) and environments with no preference information (where market-clearing lead prices offer robust guarantees). By modeling intermediate levels of information availability, it reveals how platforms can optimally incorporate their imperfect data from ratings, reviews and clicks into their ad-pricing policies to deliver the best content recommendations to customers.

Together, these five papers constitute a cohesive contribution to optimal matchmaking in two-sided marketplaces. The first two papers establish a rigorous theoretical foundation—grounded in information theory—that identifies which platform architectures achieve near-optimal performance across different preference structures. Building on this foundation, the subsequent three papers address the practical challenge of optimizing ranking and pricing mechanisms under realistic constraints, spanning from data-rich environments that enable accurate conversion-rate targeting to data-poor settings where selling customer leads or sponsored ads provide robust guarantees. I have presented these works to practitioners at Angi, Booking.com, Meta, Google, and Thumbtack, and some of the subsequent collaborations have already influenced platform policies. I continue to actively pursue these industry partnerships to further bridge the gap between theoretical insights and practical implementation.

3. Future Research Directions

Building on my theoretical foundations in marketplace design, my future research expands into two emerging areas: developing data-driven accountability tools for the home service industry and understanding how advances in artificial intelligence might transform matching markets.

3.1. Provider Quality in Home Services

Customer dissatisfaction in the home improvement industry consistently ranks among the highest across all consumer sectors, according to annual surveys by the Consumer Federation of America. In collaboration with Mohammed Alyakoob, I am assembling a comprehensive dataset linking licensed contractors in California to building permits, inspection outcomes, and online reviews. We are exploring whether permit and inspection data can improve estimates of contractor quality beyond what reviews alone provide. If successful, we plan to partner with local governments to develop objective report cards of contractor performance—analogous to Medicare.gov's quality ratings for hospitals, but applied to a sector with far less existing oversight. This research is particularly timely given California's critical rebuilding challenges following recent devastating wildfires. Tools that help homeowners evaluate contractors are essential for supporting equitable and efficient recovery efforts.

3.2. AI and Market Design

A second research direction examines how artificial intelligence may fundamentally transform matching market design and dynamics. Beyond reducing search costs, modern AI tools enable richer, more interactive preference elicitation. They allow users to describe needs in natural language and refine preferences in real time. These capabilities may transform markets where discovery frictions and quality uncertainty have historically limited efficiency. In peer-to-peer resale platforms like Facebook Marketplace, AI could help buyers identify relevant listings more quickly or assess item quality through image analysis and text interpretation. However, these tools also raise new challenges around verification, strategic manipulation, and trust. I am interested in understanding how such technologies shape incentives, information flow, and equilibrium outcomes in matching environments. I seek to anticipate unintended consequences before they become embedded in market design.

One related project will explore how AI can enhance life-critical matching systems: reducing waste in cadaveric organ allocation through better predictions of waiting times and optimized allocation rules. I have a theoretical paper on this topic (with my former PhD student Junxiong Yin) that received the Best Paper Award at WINE 2022. Junxiong's PhD thesis also ventured beyond theory with empirical analysis and simulation modeling. While this project has been on hold since Junxiong's graduation, I intend to revisit it and collaborate with transplant centers to develop AI-based prediction tools and simulation models that support clinical decision-making and organ allocation.

Across both public and private sectors, matching systems fundamentally shape how opportunity is distributed and how markets function. My research aims to understand and improve these systems through optimization theory, game theory, and empirical analysis. Looking ahead, I plan to deepen my collaborations with companies and policy organizations to bridge the gap between theoretical insight and practical implementation. My goal is to continue building rigorous yet actionable frameworks for optimizing matching markets, while contributing to real-world systems that govern access to critical goods, services, and opportunities.