Research Statement

Daniele Bracale
Department of Statistics, University of Michigan

Modern society can be seen as a network of profit-driven institutions (firms, platforms, and agencies) whose outcomes depend on each other's resources and policies. Because these decisions are interconnected, even a small policy change by one player can affect the entire network, creating uneven outcomes and economic imbalances. This highlights the need for stable equilibrium systems, where each agent maximizes its utility given others' choices, and small deviations naturally return to equilibrium. My research moves around two main goals: (1) identifying classes of systems that admit stable equilibria, and (2) designing algorithms that converge to those equilibria. I aim to build unified frameworks that connect the two economic fields of *Dynamic Pricing* and *Performative Prediction* with the statistical field of *Shape-Constrained* models.

Dynamic Pricing

Monopolistic setting. Shape-constrained statistics has a long and rich history [1, 9, 14, 21], and remains an active area of discovery. The simplest economic model where shape constraints arise, is dynamic pricing in a monopolistic setting (a single seller in the market), where the customer's valuation of an item at time t, v_t, is often assumed to follow a regression model in the item's features with an unknown cumulative distribution of the error F_0, a non-decreasing function: $v_t = \theta_0^\top x_t + \varepsilon_t$, $\varepsilon_t \sim F_0$. A transaction happens if $v_t \geq p_t$, where p_t is the price set by the seller. By observing the item feature x_t and whether the transaction has taken place or not, i.e., $y_t = 1(v_t \geq p_t)$, the seller's goal is to dynamically choose the price p_t that maximizes their revenue $E[p_t y_t]$ for $t = 1, 2, \dots, T$, where T is the time horizon. Recent works have addressed the challenge of unknown F_0 using strategies ranging from kernel methods [11] to reinforcement learning algorithms, such as bandit techniques [22] and upper confidence bounds (UCB) [18], under the assumption that F_0 satisfies Lipschitz (or stronger) conditions. In [8] we propose an isotonic regression approach to estimate F under the weaker assumption that F_0 is α-Hölder continuous for some $\alpha \in (0, 1]$, for which we derive a regret upper bound. Simulations and experiments with real-world data obtained by Welltower Inc (a major healthcare Real Estate Investment Trust) consistently demonstrate that our method attains lower empirical regret in comparison to several existing methods in the literature. Additionally, our shape-constrained approach eliminates the need for tuning parameters commonly required by existing methods.

Competitive setting. Monopolistic settings oversimplify the strategic nature of pricing, where the firms in the market actually dynamically respond to one another in a competitive environment. Although numerous sequential pricing algorithms yield low regret and converge toward Nash equilibrium, they often rely on a linear demand framework ([17]), and even approaches that incorporate nonlinear demand typically confine themselves to a specific parametric family with a predetermined nonlinear transformation ([13]), thereby reducing their applicability. Motivated by the fact that nonlinear demand models are more realistic and widely studied in various pricing applications ([12, 24, 25]), in my work [7], we (1) re-framed a widely used common assumption about the monotonicity of the *virtual valuation function* into a s-concavity assumption, strengthening the connection between the economic literature and the shape-constrained statistics; (2) proposed a semi-parametric least-squares estimation of the nonlinear demand functions, which does not require sellers to communicate their observed demand values; (3) showed that when all sellers employ my policy, their prices converge at a rate of $O(T^{-1/7})$ to the Nash equilibrium prices that sellers would reach if they were fully informed; (4) showed that each seller incurs a regret of $O(T^{5/7})$ relative to a dynamic benchmark policy. A theoretical contribution of my work is proving the *existence of equilibrium under shape-constrained demand functions using the concept of s-concavity*, and establishing regret bounds of our proposed policy. We also established new concentration results for the least squares estimator under shape constraints.

Yet, a remaining issue persists in practice: the methods in our work ([7]), as well as in common literature ([10, 15, 17]), utilize an "explore then exploit" approach, with phase lengths that depend on unknown constants ([7, 17]) and i.i.d. pricing policy in the exploration phase. In real-world dynamic pricing settings, a front-loaded exploration phase is often unrealistic and artificial. In [6], we address this by designing an algorithm that eliminates the need for explicit initial exploration phases. Moreover, we focus on a *generalized linear (single-index) demand model* setting, differently from previous works that utilize linear models ([16, 17]). In this work ([6]), we developed a UCB approach from the reinforcement learning literature. Building on this model, we propose a decentralized pricing-while-estimating policy that operates without a dedicated i.i.d. pricing design exploration phase. We prove that the individual regret scales near-optimally as $O(\sqrt{T})$ against a dynamic benchmark.

Performative Prediction

Beyond dynamic pricing, the concept of equilibrium also arises in performative settings [20]. A performative environment involves a learner (typically an institution) which goal is to learn and deploy a policy θ based on data $X \sim P$, where X represents the population features. Once the policy is deployed, the data distribution changes to $P_\theta \neq P$. This situation often appears in social systems: for example, when an institution estimates the poverty level of a population and allocates resources accordingly, the intervention may cause the population to grow and become wealthier, thereby altering the underlying data distribution. From a probabilistic viewpoint, performativity appears as a shift in distribution, i.e., the map D_θ that satisfies $D_\theta(X) \sim P_\theta$ for $X \sim P$, as agents adapt to prediction models by changing their characteristics over time. The learner's task is to estimate this distributional shift using data observed before and after deployment, and then design a policy that converges to an equilibrium θ. One such equilibrium is the performatively stable point, where deploying the policy no longer alters the data distribution once individuals respond to it, i.e., a θ_0 such that $P_{\theta_0} = P$.

- In [5], we consider a continuous-data setting (X is a continuous random variable) and adopt tools from Optimal Transport [23] to estimate the distributional map D_θ. This map can be viewed as a function that transports the pre-deployment data $X \sim P$ into the post-deployment data $D_\theta(X) \sim P_\theta$, with a cost determined by the population's utility. We provide a rate of convergence for this proposed estimate and assess its quality through empirical demonstrations on a credit scoring dataset.
- In [4], we consider the discrete setting case (X is a discrete random variable) and we introduce an ad hoc alignment procedure to consistently estimate the distribution map. Furthermore, that work proposes a policy design that enables the learner not only to recover the distribution shift but also to converge to an optimal θ.

Other Research: Neural Networks

In my career before my Ph.D. I have dedicated part of my Master's to establishing theoretical results for functional Neural Networks. In [3], we consider fully-connected feed-forward deep neural networks where weights and biases are independent and identically distributed according to Gaussian distributions. Extending previous results [19] we adopt a function-space perspective, i.e., we look at neural networks as infinite-dimensional random elements on the input space. Under suitable assumptions on the activation function, we show that: (1) a network defines a continuous stochastic process on the input space; (2) a network with re-scaled weights converges weakly to a continuous Gaussian Process in the large-width limit; (3) the limiting Gaussian Process has almost surely locally γ-Hölder continuous paths, for $0 < \gamma < 1$.

While numerous theoretical refinements have been proposed in recent years, the connection between NNs and GPs relies on two critical distributional assumptions on the NN's parameters: i) finite variance ii) independent and identically distributed (iid). In [2] we consider the problem of removing assumption i) in the context of deep feed-forward convolutional NNs. We show that the infinite-channel limit of a deep feed-forward convolutional NNs, under suitable scaling, is a stochastic process with multivariate stable finite-dimensional distributions, and we give an explicit recursion over the layers for their parameters.

Future Research Plan

As the complexity of data sets and systems increases, so too does the need for computationally efficient statistical methodologies. I will continue to develop such methodologies and algorithms to fit the real-world needs. Currently, I am working on the following topics, and I plan to pursue them in the next few years:

- *Privacy-aware signaling and equilibria.* Together with my advisor Moulinath Banerjee and the associate professor Yuekai Sun, I am studying a signaling/communication game where a platform commits to a privacy level ε (e.g., local differential privacy), and users decide how much information to reveal after observing ε. The platform and the user both want accurate recommendations, but the user also faces a privacy cost, so the game is not zero-sum. I aim to identify conditions under which the equilibrium privacy level is strictly positive (i.e., at equilibrium, $\varepsilon^* > 0$), providing design guidance for when firms should offer nontrivial privacy guarantees.

- *Equilibrium design with capacity constraints.* I study a stylized market with N sellers who each post a price, and a buyer who purchases from a single seller according to a discrete choice model. Each seller faces an initial inventory constraint and aims to sell their stock over a fixed horizon T. My goals are:
  1. *Characterize when a stable outcome (price equilibrium) exists across sellers, and is unique.* Using mild shape constraints on the distribution of the probability choice (e.g., log-concavity), I guarantee the existence and uniqueness of a Nash equilibrium.
  2. *Design learning dynamics that converge to equilibrium.* I aim to develop decentralized learning algorithms that sellers can run independently to maximize revenue under inventory constraints, balancing the trade-off between setting prices high enough to generate revenue and low enough to maintain demand and clear stock within the horizon.
- *Senior housing: data-driven pricing and access.* I am collaborating with Welltower (a leading senior housing company) to develop simple, reliable demand models and pricing methods that support stable market outcomes. The central objective is to model and predict occupancy rates for senior housing properties across the United States using property-level variables such as average rent and local demographic characteristics (e.g., age, sex, income, education). A key challenge is price endogeneity, which arises naturally in this setting. To address it, I study and evaluate plausible instrumental variables to enable more credible identification and less biased estimation of demand parameters.
- *Model-assisted transfer learning for nonparametric regression.* Together with Subha Maity (Assistant Professor at the University of Waterloo), I will develop a local-minimax framework that adaptively borrows strength from a small library of pre-trained predictors $\{f_k\}$ to estimate $\mu(x) = E[Y \mid X = x]$. The estimator selects (and softly corrects) the best local model using kernel weighting and ℓ_1 shrinkage. This yields practical procedures that (i) provably beat standard nonparametric rates when a good local surrogate exists and (ii) remain rate-optimal otherwise.

References

[1] Fadoua Balabdaoui, Cécile Durot, and Hanna Jankowski. Least squares estimation in the monotone single index model. *Bernoulli*, 25(4B):3276–3310, 2019.

[2] Daniele Bracale, Stefano Favaro, Sandra Fortini, and Stefano Peluchetti. Infinite-channel deep stable convolutional neural networks. *arXiv preprint arXiv:2102.03739*, 2021.

[3] Daniele Bracale, Stefano Favaro, Sandra Fortini, and Stefano Peluchetti. Large-width functional asymptotics for deep gaussian neural networks. *arXiv preprint arXiv:2102.10307*, 2021.

[4] Daniele Bracale, Subha Maity, Moulinath Banerjee, and Yuekai Sun. Learning the distribution map in reverse causal performative prediction. *arXiv preprint arXiv:2405.15172*, 2024.

[5] Daniele Bracale, Subha Maity, Felipe Maia Polo, Seamus Somerstep, Moulinath Banerjee, and Yuekai Sun. Microfoundation inference for strategic prediction. *arXiv preprint arXiv:2411.08998*, 2024.

[6] Daniele Bracale, Moulinath Banerjee, Cong Shi, and Yuekai Sun. Online price competition under generalized linear demands. *under review at AISTATS*, 2025.

[7] Daniele Bracale, Moulinath Banerjee, Cong Shi, and Yuekai Sun. Revenue maximization under sequential price competition via the estimation of s-concave demand functions. *arXiv preprint arXiv:2503.16737*, 2025.

[8] Daniele Bracale, Moulinath Banerjee, Yuekai Sun, Kevin Stoll, and Salam Turki. Dynamic pricing in the linear valuation model using shape constraints. *arXiv preprint arXiv:2502.05776*, 2025.

[9] Lutz Dümbgen, Sandra Freitag, and Geurt Jongbloed. Consistency of concave regression with an application to current-status data. *Mathematical methods of statistics*, 13:69–81, 2004.

[10] J Fan and I Gijbels. Local polynomial modelling and its applications. monographs on statistics and applied probability series, 1996.

[11] Jianqing Fan, Yongyi Guo, and Mengxin Yu. Policy optimization using semiparametric models for dynamic pricing. *Journal of the American Statistical Association*, 119(545):552–564, 2024.

[12] Guillermo Gallego, Woonghee Tim Huh, Wanmo Kang, and Robert Phillips. Price competition with the attraction demand model: Existence of unique equilibrium and its stability. *Manufacturing & Service Operations Management*, 8(4):359–375, 2006.

[13] Vineet Goyal, Shukai Li, and Sanjay Mehrotra. Learning to price under competition for multinomial logit demand. *Available at SSRN 4572453*, 2023.

[14] Piet Groeneboom, Geurt Jongbloed, and Jon A Wellner. Estimation of a convex function: characterizations and asymptotic theory. *The Annals of Statistics*, 29(6):1653–1698, 2001.

[15] Adel Javanmard and Hamid Nazerzadeh. Dynamic pricing in high-dimensions. *Journal of Machine Learning Research*, 20(9):1–49, 2019.

[16] Alan P Kirman. Learning by firms about demand conditions. In *Adaptive economic models*, pages 137–156. Elsevier, 1975.

[17] Shukai Li, Cong Shi, and Sanjay Mehrotra. LEGO: Optimal online learning under sequential price competition. *Major Revision at Operations Research. Available at SSRN 4803002*, 2024.

[18] Yiyun Luo, Will Wei Sun, and Yufeng Liu. Contextual dynamic pricing with unknown noise: Explore-then-ucb strategy and improved regrets. *Advances in Neural Information Processing Systems*, 35:37445–37457, 2022.

[19] Alexander G de G Matthews, Mark Rowland, Jiri Hron, Richard E Turner, and Zoubin Ghahramani. Gaussian process behaviour in wide deep neural networks. *arXiv preprint arXiv:1804.11271*, 2018.

[20] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *International Conference on Machine Learning*, pages 7599–7609. PMLR, 2020.

[21] Tim Robertson, Richard Dykstra, and FT Wright. Order restricted statistical inference. *(No Title)*, 1988.

[22] Matilde Tullii, Solenne Gaucher, Nadav Merlis, and Vianney Perchet. Improved algorithms for contextual dynamic pricing. *Advances in Neural Information Processing Systems*, 37:126088–126117, 2024.

[23] Cédric Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2021.

[24] Yi Wan, Tom Kober, and Martin Densing. Nonlinear inverse demand curves in electricity market modeling. *Energy Economics*, 107:105809, 2022.

[25] Yi Wan, Tom Kober, Tilman Schildhauer, Thomas J Schmidt, Russell McKenna, and Martin Densing. Conditions for profitable operation of p2x energy hubs to meet local demand with energy market access. *Advances in Applied Energy*, 10:100127, 2023.