Research and Teaching Statement

Kevin R. Williams

I work in industrial organization, focusing on sectors where firms operate amidst demand and competitive uncertainty and sell to customers with disparate tastes and willingness to pay. I develop new models and methods to examine the complex product offerings and pricing strategies employed by firms in these settings. To assemble my data sets, I use innovative data collection methods and leverage firm research partnerships that I have successfully cultivated over time. My research examines markets where dynamics play a pivotal role, such as airlines, and spatial settings, such as retailers. Its aim is to deepen our understanding of these important markets and to derive clear and intuitive implications for policy.

1. Dynamics and Airlines

In settings where consumers make advance purchases and supply is capacity constrained (airlines, hotels, entertainment tickets, etc.), firms' pricing serves to manage revenue in response to changing opportunity costs driven by scarcity. In my solo-authored paper "The Welfare Effects of Dynamic Pricing: Evidence from Airline Markets" (**Econometrica; 1**), I examine the resulting dynamic pricing, which requires both methodological contributions and the acquisition of highly detailed data. I analyze two factors driving airline price fluctuations. First, airlines face uncertainty about demand and dynamically update prices as bookings are realized. Such price adjustments may be welfare-improving, because they reduce the possibility of empty seats at departure. Second, prices also reflect intertemporal price discrimination that allows airlines to extract surplus, with ambiguous consumer welfare implications: Early-booking leisure travelers are more price-sensitive, whereas business travelers purchase closer to departure. To quantify these two forces, I develop a model of dynamic pricing where a forward-looking monopolist airline adjusts prices based on demand shock realizations and changes in overall consumer willingness to pay over time. I exploit the firm's pricing decisions to separately identify consumers' preferences versus the magnitude of demand uncertainty, and estimate the model with novel data on prices and bookings I derived from monitoring changes in more than 700,000 airline seat maps. My central finding is that flexible price adjustments are on net welfare-enhancing. Firms and leisure consumers benefit from dynamic pricing. However, although dynamic pricing helps ensure seat availability for business consumers, these travelers are made worse off due to higher prices closer to flight departure. As one of the 10 most-cited papers published in *Econometrica* between 2022 and 2023, this paper appears on graduate reading lists at universities such as the University of Chicago and University of Pennsylvania.

While profit-maximizing firms are a natural benchmark, I have also established unique insights by looking at decision-making processes within the firm. In "Organizational Structure and Pricing: Evidence from a Large U.S. Airline" (**Quarterly Journal of Economics; 2,** with Hortaçsu, Natan, Parsley, and Schwieg), I demonstrate that relaxing the common assumption that firms are unitary profit-maximizing decision-makers is crucial for understanding firms' actual pricing and welfare. I exploit comprehensive data and internal models at a large U.S. airline to establish that pricing is subject to multiple biases. These include relying on a heuristic that does not internalize product substitution of any kind, using persistently biased demand forecasts, and pricing on the inelastic side of internally estimated demand curves. Standard economic models would conclude that this is suboptimal. However, I show that observed prices can be rationalized by accounting for organizational structure under which multiple departments make separate contributions to pricing. These decisions include the fare menus assigned to itineraries, demand forecasts, and capacity decisions. I model a non-unitary firm where departments submit inputs to the observed pricing heuristic and establish that given other departments' heuristic inputs, providing biased inputs is revenue-maximizing for the firm. For example, decreasing the demand forecast from its observed biased levels does not increase overall revenues, assuming other departmental inputs that contribute to the pricing heuristic remain constant. Therefore, department managers choose inputs in a boundedly rational way given their delegated decision rights. There are substantial welfare consequences from prices being set through delegated decision-making. Simulating the prices the firm would charge if it were a unitary decision-maker results in substantially lower welfare due to higher prices than I estimate under observed practices. These findings have implications for modeling supply broadly, because firms in a wide variety of sectors use the observed organizational structure.

While papers (1) and (2) analyze monopoly markets, "Dynamic Price Competition with Capacity Constraints" (**submitted; 3**, with Betancourt, Hortaçsu, and Öry) examines dynamic pricing in the common but theoretically more challenging setting of competition. Firms consider not only their own remaining capacity, as in (1) and (2), but also their rivals' remaining capacities, because this impacts future competition. I provide sufficient conditions for the existence and uniqueness of pure-strategy Markov perfect equilibria in discrete time. I use the continuous-time limit to highlight new economic forces and show that little intuition from monopoly dynamic pricing models carries over to the oligopoly case. For example, prices can be lower if firms have lower initial capacities. This finding runs counter to the intuition that less inventory results in higher prices. Moreover, firms' strategies encourage the firm with the lowest capacity to sell out first, whereas a monopolist benefits from preserving product

variety over time. The way competition promotes scarcity can exacerbate misallocation by causing firms to sell capacity too quickly, and this in turn harms both consumers and firms. I formalize this new welfare effect, which I call the *Bertrand scarcity trap*, and contrast market outcomes with alternative pricing mechanisms in simulation studies. I demonstrate that pricing algorithms can increase consumer surplus and revenues relative to the competitive equilibrium outcome.

In "Intertemporal Price Discrimination in Sequential Quantity-Price Games" (**Marketing Science; 4**, with Dana), I study how dynamic price competition is affected if firms can choose capacity first. Using a two-period model without demand uncertainty and homogenous products, I find that strong competitive forces can prevent firms from using intertemporal price discrimination even if consumers' willingness to pay increases over time. Intuitively, multiple sales periods create a costless arbitrage opportunity for capacity-constrained firms: Shifting early, price-sensitive consumers to a rival increases a firm's market share in the second period, when price-insensitive consumers make their purchasing decisions. However, rival firms have similar strategic incentives, leading to a uniform-price equilibrium. I demonstrate that intertemporal price discrimination is feasible only if firms implement per-period inventory controls, which prevent firms from selling too much of their capacity at a given price. Therefore, inventory controls can soften price competition, in addition to their established role in managing demand uncertainty.

In ongoing research with Aniko Öry, I explore how competitive interactions are shaped by firms' reliance on simple heuristics when faced with complex decisions. Building on the findings in (**2**), I introduce the concept of heuristic equilibrium, where firms use simplified and sometimes imperfect models, such as ignoring substitution patterns or not updating beliefs about rival states based on observed actions. I establish conditions for the existence of pure-strategy heuristic equilibria and highlight the alignment between my framework and airline pricing practices using detailed data from competing U.S. airlines. Market outcomes assuming heuristic equilibria differ substantially from market outcomes assuming Markov perfect equilibria as characterized in (**3**). This finding has important implications for competition policy and the broader modeling of competitive interactions, particularly in the era of algorithmic pricing.

While my work on airlines focuses on how firms use prices to manage demand uncertainty, in "Aiming for the Goal: Contribution Dynamics of Crowdfunding" (**American Economic Review; 5,** with Deb and Öry), I explore another approach: crowdfunding. I analyze a contribution game where randomly arriving buyers decide whether to pledge to purchase a product before its launch, which happens only if enough buyers commit before a deadline. A

representative, long-lived donor values a successful launch, but their budget to contribute over time is their private information. Buyers face a coordination problem, which I show is alleviated by the donor signaling their budget or private willingness to contribute over time. I derive bounds on the effect of dynamic signaling on the probability of project success and establish that making the donor's budget public information would result in the lowest probability of success. I highlight the trade-offs between maximizing the donor's and buyers' payoffs and explore various contexts in which my model can be applied more broadly. I then focus empirically on the case of Kickstarter. Using new data that distinguish pledges for products from donations at 12-hour frequencies, I provide descriptive evidence consistent with equilibrium predictions and calibrate the model to quantify how dynamic signaling facilitates buyer coordination. I find a significant gap in the success rate between the success-maximizing and the success-minimizing equilibria and show that alternative platform policies result in even lower success rates.

2. Spatial, Retail Markets, and Demand Methodologies

My second research area examines spatial variation in the retail sector, where price discrimination is also widely employed. In "Zone Pricing in Retail Oligopoly" (**American Economic Journal: Microeconomics; 6,** with Adams), I show that large brick-and-mortar home improvement retailers charge different prices based on the geographic location of their stores but only to a limited extent. This form of market segmentation may seem surprising given that costs, demand, and competition in a city store are very different from those in a small-town store. I compile price and quantity data by monitoring stores' online reports of inventory levels and estimate a model of zone pricing. I then simulate counterfactuals involving finer price discrimination (market-level pricing) and uniform pricing, finding that Lowe's prefers coarser pricing, whereas Home Depot benefits under finer market segmentation. Zone pricing shields rural consumers from high prices where firms might otherwise exercise their market power at the expense of higher prices in competitive, urban markets. Thus, overall, zone pricing benefits consumers.

While the previous paper focused on spatial variation in prices, in "Product Variety, Across-Market Demand Heterogeneity, and the Value of Online Retail" (**RAND; 7,** with Quan), I explore spatial variation in product demand and availability. I estimate the consumer surplus gains created by e-commerce using all shoe transactions from the e-commerce retailer Zappos.com. Methodologically, I address the challenge that most shoes do not sell in each location, any given week (for example, sandals rarely sell in Alaska in any given week, and less in Minnesota in the winter). I show how random-coefficient demand models can accommodate zero-sale observations, where existing approaches fail because the log of zero is

undefined. My estimates suggest that accounting for heterogeneity in preferences across geographies and the fact that local retailers target product assortments to the tastes of local consumers, previous estimates on the welfare gains from e-commerce may be biased by upwards of 75%.

I have two related methodological papers that arose from studying markets where product sales are low. In "Extracting Characteristics from Product Images and its Application to Demand Estimation" (**R&R RAND; 8,** with Quan) I leverage machine learning techniques and product images to estimate substitution patterns that are otherwise difficult to quantify due to this sparsity in sales. In "Demand Estimation with Infrequent Purchases and Small Market Sizes" (**Quantitative Economics; 9,** with Hortaçsu, Natan, Parsley, and Schwieg), I present an alternative demand methodology that allows for zero-sale observations. I augment standard random coefficient demand models with a Poisson arrival process and use a hybrid-Gibbs Bayesian estimation approach. I demonstrate that the estimator performs well in situations where alternative methods produce biased demand estimates. I also demonstrate how to use the estimator in an empirical application to airlines, where search counts are derived from clickstream data.

In my final line of research, I analyze consumer movement using novel geolocation data. Distance to the consumer is a key strategic choice in retail; yet, most work ignores that travel costs are endogenous to the retailer's location decision. In "Distributional Impacts of the Changing Retail Landscape" (**working paper; 10,** with Cao, Chevalier, Handbury, and Parsley), I propose a novel ring instrument using spatial demographic data to address distance endogeneity and show that it results in substantially higher consumer travel costs than assuming distance is exogenous. Using an instrumental variables approach, I estimate preferences across income quartiles and geographies for general-merchandise stores based on smartphone geolocation data and quantify the welfare effects of the so-called retail apocalypse, or the mass closings of retail establishments. As I show in other work, smartphone data are aptly suited for studying spatial retail questions, because these data cover broad sections of the U.S. population and exhibit movement patterns similar to conventional survey data across income deciles ("Measuring Movement and Social Contact with Smartphone Data: A Real-Time Application to COVID-19," **Journal of Urban Economics; 11,** with Couture, Dingel, Green, and Handbury). In **10**, I demonstrate substantial differences across retail chains in their contributions to welfare for different demographic groups and estimate that welfare per store trip has not substantially declined over the past decade.

The fact that stores strategically locate based on their targeted demographics suggests that leveraging these entry decisions can have beneficial impacts for public health initiatives. In

"Distributional Impacts of Retail Vaccine Availability" (**Journal of Urban Economics; 12,** with Chevalier, Schwartz, and Su), I show that including dollar stores in the Federal Retail Pharmacy Program would substantially decrease the distance to vaccine sites for low-income, rural, and minority U.S. households.

Teaching Statement

In my teaching, I aim to strike a balance between theoretical rigor and practical application by stressing general principles that are real-world and industry relevant. Most of my teaching has been in Competitive Strategy, one of the most popular electives at Yale SOM. This course explores strategic decisions within an organization and uses modern game theory to showcase how firm profitability depends on the actions of other firms. I update and improve the course every year. For example, in 2023 I added a class on competition in data that explores how data can create market entry barriers, which I tie into privacy and regulation. In 2024, I added a class on algorithmic pricing and its role in softening price competition, with applications from gasoline and rental housing markets. Student feedback has been positive, with evaluations consistently ranging from 4.0/5 to 4.8/5. In 2023, I was voted Favorite Professor by students in the Healthcare Management program. I have also taught Competitor, the second required Microeconomics course, and Basics of Microeconomics in the executive MBA program.

Outside the classroom, I have mentored nine predoctoral fellows during my tenure at Yale. I have purposefully focused on supporting underrepresented minority groups in higher education with the goal of encouraging them to pursue a PhD. These exceptionally talented research associates have been instrumental in advancing data- and compute-intensive airline research. Hayden Parsley is a coauthor of (2), (9), and (10). Jose Betancourt is a coauthor of (3). These predoctoral students have been placed into top PhD programs including those at Brown, Columbia, Duke, and Yale. I also mentor three first-year college students at Grace Hopper College every year. I regularly attend the PhD IO workshop to provide students with feedback, have supervised an undergraduate senior thesis writer in economics, and mentor high school students interested in pursuing economics.

Updated: July 25, 2024