# Keith Marzilli Ericson

Personal Statement May 2022

## Research Since My Promotion to Tenured Associate Professor

My research examines healthcare markets and behavioral economics and has been published not only in top economics journals, but also in *Psychological Science* (the highest ranked empirical journal in psychology) and *Management Science* (the flagship management journal). My work has received media coverage in various outlets, including the *New York Times*, *Washington Post,* and *Harvard Business Review*.

Healthcare comprises nearly a fifth of US GDP, and the health sector is often viewed as inefficient. My research is devoted to understanding the organization of healthcare markets and how to improve them. One challenge in the healthcare sector is the management of firms delivering healthcare services, and a line of my recent research examines how teams and organizational boundaries affect healthcare productivity (Agha et al. *Management Science* forthcoming, *AEJ Policy* forthcoming, Ericson et al "Dangerous Prescribing…"). Another challenge is the design of health insurance markets which finance much of healthcare spending, and a line of my research examines insurance markets and consumer choice, showing providers' negotiated prices vary substantially across insurers (Craig, Ericson and Starc 2021 *Journal of Health Economics*), and evaluating interventions to improve health insurance choices (Ericson et al 2017 *Health Affairs*, Ericson et al "Nudging Take-up of Subsidized Insurance…").

Healthcare and insurance research requires understanding individual decisions and human psychology. Insurance aims to reduce risk, so understanding individual attitudes toward risk is important to design insurance contracts well and to evaluate individuals' decisions. I have shown how to estimate risk aversion from choices in insurance markets even when consumers have incorrect beliefs (Ericson et al. 2021 *Economic Journal*), and shown that modeling liquidity-constrained individuals can explain puzzling behavior in insurance markets (Ericson and Sydnor "Liquidity Constraints…"). I have also shown how market equilibrium is affected by defaults, and how to optimally choose defaults when inertial behavior may result from either real switching costs or choice frictions (2020 *Games and Economic Behavior*).

Finally, an important part of my research agenda has been to shape how economics approaches core problems. In a pair of influential papers, I examine time preference and intertemporal choice. In Cohen et al (2020 *Journal of Economic Literature*), I argue the most commonly used tool to measure time preference does not measure what many think it does, and identify how better experimental paradigms yield different results. In Ericson and Laibson (2019 *Handbook of Behavioral Economics*), I identify a series of unanswered questions in intertemporal choice. Finally, while more choice is often thought to make consumers better off, market failures in insurance can lead choice to backfire. In Ericson and Sydnor (2017 *Journal of Economic Perspectives*), I examine the potential gains from having a choice of health insurance coverage levels and how those gains are affected by market regulations and individual mistakes.

## Managing Healthcare Organizations

Healthcare is delivered by diverse teams of providers, often working across organizational boundaries to solve complicated problems. Coordinating complex tasks and sharing information within teams is difficult. Organizational boundaries can affect the flow of information and the incentives that team members face. In a series of papers, I examine how organizational boundaries affect healthcare productivity and quality— with implications for evaluating competition policy— and develop new measures of teams in organizations.

In Agha et al (forthcoming *Management Science*), I examine the teams constructed by primary care physicians ("PCPs") when they make referrals to specialists. **I show that healthcare utilization drops without any detectable loss of quality when PCPs concentrate their referrals to a small number of specialists rather than diversifying their referrals.** Establishing this point required overcoming two main challenges. The first was to develop a measure of the extent to which PCPs concentrated their referrals. The second was to identify a causal effect. Previous research on referral teams only looked at how many physicians were involved in care delivery. My new measure— referral concentration— introduced a new dimension: repeated interactions between team members. Measuring this for each PCP in Medicare claims data and Massachusetts commercial claims data, I find that referral concentration varies widely across PCPs, even among PCPs in the same organization.

I use two identification strategies to estimate the causal effect of referral concentration on healthcare. One strategy uses both PCP fixed effects and fixed effects for the particular specialists seen, so results are identified, for instance, by comparing two patients with the same conditions seeing the same cardiologist but coming from primary care physicians that vary in how closely they work with that cardiologist. Chronically ill patients treated by PCPs with a one standard deviation higher team referral concentration have 4% lower healthcare utilization on average, with no discernible reduction in quality. I show similar effects using a difference-in-differences strategy, in which I examine what happens when a patient must switch between PCPs with different levels of referral concentration due to a move across regions. These results show that managers of healthcare firms can increase healthcare efficiency by directing referrals to create repeated interactions. The results can also inform the design of health insurers' networks, which can steer patients, and policy initiatives, such as Accountable Care Organizations.

In Agha, Ericson, and Zhao (forthcoming *AEJ Policy*), **I examine how the boundaries of the firm affect healthcare efficiency.** Transaction costs and imperfect information can make it difficult to coordinate production across firm boundaries. I show that healthcare costs are lower without any worsening of outcomes when the providers treating a patient work within the same organization. This work needed both a way to measure the extent to which a patient received healthcare in the same firm and a way to measure a causal effect. I introduce the concept of *organizational concentration*, which measures the distribution of a patient's healthcare encounters across organizations. For causal identification, I use two identification strategies. First, I examine what happens when patients move to regions where other patients typically concentrate their outpatient visits in a small set of firms. I find that the healthcare utilization of these movers falls, even holding constant traditional measures of healthcare fragmentation. Second, I examine patients whose primary care provider (PCP) exits the local market.

Since patients may endogenously sort to new PCPs based on changes in their health status, I use an instrumental variable strategy that leverages mean reversion to predict the change in a patient's assigned PCP's average organizational concentration. **I find that switching to a PCP with 1 standard deviation higher organizational concentration reduces utilization by 21%, again with no discernable reduction in quality (and perhaps even quality improvements for diabetes care)**.

**I also find that the concentration of care within an organization is more important than care fragmentation across providers.** This result has important managerial and policy implications that run counter to conventional wisdom. While previous work has emphasized the economic losses from concentrating providers into fewer organizations, I find evidence for a potential gain from integrating providers into the same organization. Since these efficiency gains come from reduced coordination costs for multiple providers treating the same patient, these results help identify types of mergers that are more likely to be efficient. Mergers that simply bring together competing providers who rarely would have treated the same patient will not have the efficiencies identified in this paper.

In Ericson, Sacarny, and Zhou ("Dangerous Prescribing…"), **I investigated how fragmented healthcare delivery is connected to dangerous prescribing, particularly of opioids.** Deaths from opioids result not only from excessive opioid prescribing, but frequently from prescription drugs that have dangerous interactions with opioids. Lack of coordination is a potentially important cause of dangerous prescribing and a large, disorganized care team could inadvertently prescribe medications with dangerous opioid interactions. These concerns have motivated policy interventions such as prescription drug management programs.

However, contrary to this prediction, I show that dangerous opioid prescribing is lower in regions that have more healthcare fragmentation (measured in a variety of ways). This is surprising, since previous research has shown that more fragmented regions use more healthcare and do not have better outcomes. However, this cross-regional correlation is not a causal relationship. Examining the issue with an identification strategy that focuses on patients who move across areas with differing levels of fragmented care delivery, **I find a precisely estimated zero effect of fragmentation on dangerous opioid prescribing and most measures of opioid prescribing volume.** These results indicate that healthcare fragmentation is not driving regional variation in dangerous prescribing. Moreover, they indicate that using simple regional correlations about healthcare practices to make causal inferences is difficult.

## Insurance Markets

### Healthcare Prices

High healthcare costs are a major public policy concern. To investigate why prices are high, research has focused on price variation between providers, for instance comparing high- versus low-priced hospitals. If quality is similar but price is substantially different between providers, there may be opportunities to reduce healthcare costs without compromising quality.

In Craig, Ericson, and Starc (2021 *Journal of Health Economics)*, I investigate how negotiated transaction prices paid by different insurers for the same procedure at the same hospital vary. If prices are driven by

quality differences or patient demand for particular providers, we would expect to see that insurers negotiated similar prices at a given provider. **However, I find that price variation between insurers is about as important as price variation between hospitals in the average prices they charge.** (The standard deviation in average price across insurers is approximately the same as across hospitals.) Prices at the level of the insurer-provider pair are difficult to observe. This project required gaining access to data that could be used to estimate negotiated prices, access that was opposed by insurers. While many people focus on health insurance premiums, this work highlights that both consumers and self-insured employers should also pay attention to the less-visible prices that an insurer has negotiated with providers. These provider prices will affect both the out-of-pocket cost-sharing that consumers face and the total cost to self-insured employers.

I further show that even within an insurer, negotiated prices at a provider across the different contracts offered by that insurer. I develop a model of negotiation incentives, comparing fully-insured plans (the insurer is the risk-bearing entity) to self-insured plans (the employer is the risk-bearing entity). In fully-insured plans, the insurer benefits dollar-for-dollar from lower negotiated prices, while for self-insured plans, the insurer only benefits from lower negotiated prices if they lead to more employers working with the insurer. Consistent with the model's predictions, **I find that a plan for only self-insured enrollees would pay about 8% more—at the same hospital and with the same insurer— than a plan only for fully insured enrollees.**

## Interventions to Improve Insurance Choice

**Defaults.** Default rules determine what happens to individuals who are inattentive and do not make a decision. In Ericson (2020 GEB "*When Consumers Do Not …*"), I developed theory showing **how automatic renewal default rules affect market equilibrium and how to set them optimally.** Markets vary in whether the default option is to automatically renew the purchase of the product or service (most subscription services), to cancel the subscription (e.g. some states' health insurance exchanges), or to automatically switch to an alternative, cheaper product (e.g. Medicare Part D's program for low-income individuals). Firms respond to individual behavior when setting prices, and so respond to defaults. I showed how the choice of default affects the price elasticity of repeat demand, and thus the pricing pattern of firms.

**Default rules, in effect, have externalities, and the privately optimal default rule will not typically coincide with the socially optimal default rule.** For instance**,** an automatic switching default can lead to fewer individuals actually switching in equilibrium by increasing the price pressure on firms and lowering the price dispersion in the market. An automatic switching rule, like that used by Medicare, builds in an institution that acts like a savvy shopper. **I then showed how the socially optimal default can be determined from simple choice experiments.**

**Nudges.** Insurance markets are unlikely to work well if consumers don't shop well for plans, and there is policy interest in helping consumers make better choices. One concern is inertia: when prices change, sometimes dramatically, many consumers do not reoptimize even though doing so could save them money. In Ericson et al. (2017 *Health Affairs*), **I conducted a randomized intervention to encourage enrollees in an ACA Marketplace to shop for plans.** I tested the impact of personalized information

about the savings on insurance premiums that they could get if they switched plans, and generic communications that simply emphasized the possibility of saving. The personalized and generic messages both increased shopping on the Marketplace's website by 23 percent, but neither type of message had a significant effect on plan switching. These results indicate that inertia is not simply the result of inattention. Consumers may still strongly prefer their current plans (perhaps due to real switching costs), or may face barriers to improving choices even after they began paying attention (such as an inability to determine which plan was better for them or an inability to follow through with switching intentions due to procrastination or forgetting).

In work in progress (Ericson et al. "Nudging Take-up of Subsidized Insurance…"), **I evaluate the impact of an intervention I conducted with an ACA Marketplace to increase take-up of free or low-cost health insurance.** I examine the impact of removing significant non-price barriers (information frictions, hassle costs) and price barriers on take-up of insurance, and see what types of individuals are screened out by price versus non-price barriers. I worked with policy-makers on the Massachusetts Health Connector to develop a set of low-cost interventions to increase enrollment. In collaboration with the Connector, we created a simplified enrollment pathway that reduced administrative barriers. I find that the most intensive intervention increases enrollment by about 10% over baseline, and that there is important heterogeneity across groups in the intervention that is most effective. In ongoing work, I am linking the results with healthcare claims data to model the impact of this intervention on market equilibrium and developing a theoretical model to distinguish among behavioral mechanisms by which our intervention affected choice and conduct welfare evaluation.

### Risk Adjustment and Insurance Markets

Many health insurance markets use risk adjustment to facilitate healthcare market competition. Risk adjustment schemes transfer money from insurers who enroll healthier patients to those who enroll sicker, more expensive patients. The goal is to reduce distortions in pricing and in contract design due to adverse selection.

While risk adjustment typically aims to predict the annual costs of individuals enrolled in an insurance plan for a full year, partial-year enrollment is common due to churn in the markets. Risk adjustment is particularly difficult when enrollees church quickly from one plan to another. In Ericson, Geissler, and Lubin (2018 *AJHE),* I find that existing risk adjustment schemes in ACA Marketplaces and state Medicaid programs **will systematically undercompensate firms for partial-year enrollees**: due to missed diagnoses, partial-year enrollees for whom no diagnosis is observed are sicker on average than full-year enrollees for whom no diagnosis is observed. **I develop a model of risk adjustment for partial-year enrollment, show that existing ad hoc solutions do not solve the problem, and show how to scale up payments for partial-year enrollees' observed diagnoses to give firms the correct incentives.**

Giving people choice over health insurance can create value when people sort based on preference heterogeneity, but it also creates challenges related to both consumer confusion and adverse selection. In Ericson and Sydnor (2017 *Journal of Economic Perspectives*), I provide a framework to assess the tradeoffs involved in giving consumers options of levels of health insurance generosity. I show that institutions, such as risk adjustment, that reduce how adverse selection affects plan prices can make

offering choice over coverage levels beneficial to fully informed consumers. Yet quantitatively, the gains from having choice are small and eroded quickly with imperfect risk adjustment or confused consumers.

### Work In Progress: Paying for Prescription Drugs

In recent work, I am examining how to design prescription drug payment policy. The key tradeoff in paying for prescription drugs is incentives for innovation (suggesting high prices) versus efficiency gains from using drugs when they are needed (suggesting lower prices). In pre-tenure research (Cutler and Ericson 2010 *PharmacoEconomics*), I showed that because pharmaceutical prices have high markups over marginal cost due to patent protection, accounting for the social cost of pharmaceuticals changed the recommendations from cost-effectiveness analysis on how many people should take a drug.

My current research is funded by a grant from the National Institute of Healthcare Management and examines payment policy for Physician-administered drugs in Medicare Part B. I aim to address both an empirical question—how does current payment policy affect the price path of physician-administered drugs over time?—and a normative question—how can payment policy be improved? A key insight is a single price per unit of drug cannot efficiently set both incentives for use of the drugs and incentives for innovation. However, more complex contract designs, such as two-part tariffs (in which pharmaceutical firms get a fixed fee plus an additional payment per unit sold), can set the right incentives on both margins.

## Individual Decision Making

### Risk Preferences and Insurance

When a consumer has a high willingness to pay for insurance, that can result either from a high anticipated probability of making an insurance claim ("risk perceptions") or from a strong desire to avoid risk ("risk preferences"). Separating these two sources of demand is important for predicting how policies will affect insurance markets and evaluating their welfare consequences. However, doing so is difficult because we do not typically observe individuals' risk perceptions. Researchers often assume that consumers have rational expectations about the probability they will make a claim and use claims data to identify these risk perceptions. However, a large literature in behavioral economics suggests that individuals do not accurately estimate probabilities.

In Ericson, Kircher, Spinnewijn, and Starc (2021 *Economic Journal*), **I show how to identify consumers' risk preferences even if they have incorrect beliefs about their healthcare usage**. **The innovation in this paper was to show how to separately identify risk aversion from risk perceptions using only insurance choice data by exploiting variation in prices and plan characteristics.** The insight is that if demand is driven by high risk aversion instead of high risk perceptions, willingness to pay for more insurance will decline more quickly as the insurance contract approaches full insurance. That is, from the perspective of risk aversion, lowering the deductible from $5000 to $4900 is much more important than lowering the deductible from $100 to $0.

I then applied these theoretical results to empirically estimate bounds on the distribution of risk preferences and risk perceptions for enrollees in the Massachusetts health insurance exchange. This methodology can be used to learn about preferences and perceptions when claims data is not available, as well as used to evaluate the accuracy of risk perceptions when claims data is available.

In another line of work, I examine the source of risk preferences themselves. A large behavioral economics literature argues that individuals display puzzlingly risk-averse behavior. Moreover, some choices were interpreted as clear evidence of irrationality—for instance, paying $1000 in annual premiums to lower the deductible by $500. (This is a dominated choice, as if you choose the lower deductible but higher premiums, you would have less money at the end of the year regardless of whether you make a claim.)

Many people face serious liquidity constraints that limit their ability to smooth financial shocks over time by borrowing, which can affect their desired protection against risk. In Ericson and Sydnor ("Liquidity Constraints and the Value of Insurance"), I develop new insights into how individuals should behave in insurance markets if they have liquidity constraints. I build a model of insurance demand that incorporates limits to borrowing and saving, and show that **for liquidity-constrained individuals, insurance provides a consumption-smoothing benefit (smoothing expenditure across time within the year) beyond its traditional risk-protection benefit (smoothing consumption across states)**.

I use both theory and simulation to show how this model explains patterns of insurance choice that have puzzled economists, including the purchase of more insurance even when premiums are so high as to be dominated in the standard model, or the purchase of insurance for events that are certain to happen. The model gives us new insight into how insurance contracts can be designed. For instance, a classic result is that optimal insurance contracts should take the form of a deductible with full insurance thereafter (in the absence of moral hazard), but liquidity-constrained individuals benefit from insurance designs that smooth cost-sharing over time, such as coinsurance or copays.

## Time Preference

Most important decisions involve intertemporal choice: trading off costs and benefits that occur at different points in time. Do I spend money today, or save it for later? Pre-tenure, I developed a body of research that advanced our understanding of how individuals make such intertemporal choices. For instance, discounted utility is the primary framework economists have used to analyze intertemporal choice. However, in Ericson et al. (2015 *Psychological Science*), I showed that heuristic models of decision-making predict better out-of-sample than discounting models for the "money earlier or later" (MEL) task used in hundreds of economics papers and psychology papers. In recent research (Ericson and Noor "Delay Functions"), **I directly test and reject assumptions underlying the discounted utility model. The standard model of time preference—even with present bias—does not describe what many participants are doing when they make choices involving getting money or later.**

Based on my prior work in this area, I was invited to write a review of the time preference literature for the prestigious *Journal of Economic Literature* (Cohen et al. 2020). The purpose of these reviews is typically to summarize the current state of the literature and also to point out areas that researchers ought to prioritize. One of the areas for future work highlighted in my review is quite fundamental. Time preference is about preference over when utility is consumed, which is distinct from how consumption is financed. The workhorse tool for studying intertemporal choice obscures this aspect of time preferences because it focuses only on the financing problem—the would you prefer "money earlier or later" task. The difference between having the money (a financing problem) and consuming

with the money is overlooked. Measured time preferences differ substantially between consumption and financing tasks. **I develop a model showing the (implausible) assumptions necessary to interpret "money earlier or later" decisions as measuring time preferences over consumption.** Then in Ericson and Laibson (2019 *Handbook of Behavioral Economics*), I synthesize existing models of intertemporal choice, introduce the concept of "present-focused preferences" to describe the large class of models that prioritize present flows of experienced utility, and identify major open questions about intertemporal choice.

Because measuring time preferences over consumption is difficult, most studies focus on money. My work in McClure et al. (2007 *Journal of Neuroscience*) is one of the few early studies to measure time preference over actual consumption. Thirsty subjects deprived of liquids made choices of how much liquid to consume at different points in time. These data showed evidence of present bias in time preference on the time horizon of minutes**.** In more recent work (Dorison et al 2020 *PNAS*)**, I used another real consumption task to investigate time preference.** We examined decisions of how much to smoke at different points in time in the laboratory to measure the impact of emotional manipulation on time preference.

# Summary Of Research Prior to Tenure

## Health Insurance Choice and Imperfect Competition

An influential line of my work has focused on the regulation and performance of health insurance exchanges—markets in which insurance plans must take all comers at a posted price—which have been the backbone of recent American insurance market reforms, such as the Affordable Care Act and Medicare Part D.

In Ericson and Starc (2015 *Review of Economics and Statistics*), **I showed the importance of modeling imperfect competition for understanding the impact of pricing regulation**. I examined modified community rating, which limits the degree to which prices for insurance can vary across people. Under perfect competition, these regulations simply redistribute money from expensive enrollees to cheaper enrollees. However, under imperfect competition, we find that these community rating regulations reduce firm markups over cost. They thus lower firm profits, increase consumer surplus, and make the allocation of consumers across plans more efficient. In Ericson and Starc (2016 *Journal of Health Economics*), **I examined a policy that aimed to facilitate competition: standardization of health insurance plans.** Standardization is relevant not only for health insurance, but for many complex contracts, such as cell phone plans and credit cards. Using a natural experiment, I found no evidence of increased consumer price sensitivity, but large shifts in how consumers valued attributes of health insurance plans. These results imply that standardization may not achieve the goal of increasing price competition. Finally, **a pair of papers established key facts about consumer price sensitivity and valuation of provider networks in these new markets** (Ericson and Starc 2012 *American Economic Review Papers and Proceedings*, 2015 *American Economic Review Papers and Proceedings*).

Consumers purchase health insurance plans on an annual basis, but their initial decisions can have long-lasting consequences. In Ericson (2014 *American Economic Journal: Economic Policy*), **I showed how insurance firms respond to individuals' inertia.** I examined the launch of Medicare Part D, which

provides prescription drug insurance for the elderly. I first show that initial decisions persist over time. Rational firms should respond to this inertia. Typically, this takes the form of introductory offers, but these are prohibited in Part D. I showed that firms attempt to replicate introductory pricing with invest-then-harvest pricing, entering the market with low prices, acquiring market share, raising prices on consumers in later years. In those later years, new products and/or firms enter the market at low prices to appeal to consumers entering the market for the first time. I estimate that existing plans are 20% more expensive than equivalent, newly-introduced plans.

### Behavioral Economics

**Loss aversion**—the tendency of individuals to attach more weight to a loss than an equivalently sized gain—is a crucial component of Kahneman and Tversky's Nobel-prize winning Prospect Theory. Previous literature had assumed loss aversion around ownership-based reference points accounted for the endowment effect. However, in Ericson and Fuster (2011 *Quarterly Journal of Economics*), **I tested an alternative theory: that individuals' expectations about future outcomes determine the reference point.** My experiment held ownership and other factors constant and randomized the probability an individual expected to receive an item. Because of this paper's impact on the field, I wrote a review article examining the state of the literature on the endowment effect (Ericson and Fuster 2014 *Annual Review of Economics*). In it, I proposed a theory of multiple reference points that unifies the disparate results in the literature and highlights the importance of the *salience* of different potential reference points.

Answers to survey questions about sensitive topics might be biased towards social norms, which presents a challenge for research studying the LGBT population. In Coffman, Coffman, and Ericson (2017 *Management Science*), **I showed that social desirability bias leads to underestimates of the size of the LGBT population and anti-LGBT sentiment.** I used a series of experiments to compare estimates from the standard methodology of asking sensitive questions to those from a "veiled" methodology (also called the Item Count Technique) that precludes inference about an individual but provides population estimates. The veiled method increased self-reports of anti-gay sentiment, particularly in the workplace. Moreover, it leads to a substantial increase (65%) in the fraction of the population that identifies as LGBT. This work not only contributes to the literature on sexual orientation, but also to measuring sensitive behaviors more broadly: I evaluated the validity of the methodology using the first large-scale placebo test for the Item Count Technique ever conducted.

## Teaching and Advising

My teaching philosophy is informed by psychological research and relies on active learning. The form that active learning will take will vary by context. For larger classes, class time can be used to work through problems or examples that illustrate key concepts. Class time thus builds on, rather than duplicates, assigned reading by having students solve problems, make new connections, and identify topic areas that need clarification. Effective active learning not only solicits questions from students, but requires students to prepare for class by either answering brief questions in advance of class or being ready to answer questions posed to them during class.

## Business Analytics Teaching

As the Faculty Director for Undergraduate Business Analytics, I took on the task of redesigning the three-course business analytics curriculum (statistics, econometrics, optimization, and data science) required of all approximately 700 undergraduate business majors each year. This was a challenging assignment, since we have large heterogeneity in student interests and abilities, and was critical for the school's future because of the increasing importance of quantitative and statistical methods for modern management.

I redesigned the introductory required analytics course and modernized it to focus on "making decisions with data," shifting it away from null-hypothesis testing toward working with data, estimating effects, and quantifying uncertainty. I developed a set of cases and business problems to show how statistical methods can extract insights from data to inform business decisions. I pushed students not only to be able to work with the data, but also to communicate their insights. To do this, I introduced a writing assignment in which students must analyze a firm's randomized experiment and draw conclusions on whether they should roll out a new product. Even though this class was a large lecture, I used an active learning approach in which students solved problems and worked in small groups. In addition to the lecture, students had twice-weekly discussion sections led by teaching faculty, and so I also managed a team of multiple teaching faculty and a staff of undergraduate TAs.

## Behavioral Economics Teaching

I have also developed a series of classes on behavioral economics for different audiences. I developed an interdisciplinary course for undergraduates on "The Psychology of Decision-Making: Implications for Business and Public Policy," using active learning, in-class experiments, and in-depth cases. My development of this course received a grant from the Provost's program to encourage interdisciplinary teaching. Second, I developed an MBA elective on "Improving Your Decisions," applying behavioral economics and game theory to personal and business decision situations. Finally, I have developed and taught a doctoral-level course on behavioral economics.

I have advised a variety of doctoral and post-doctoral students. Though my department does not have a doctoral program, I have also been sought out as a mentor for PhD students from the Boston University Department of Economics. I have advised a series of students from the Boston University economics department: Julie Shi (Peking University), Timothy Layton (Harvard University, Dept. of Healthcare Policy), Emily Gee (U.S. Dept. of Health and Human Services), Wenjia Zhu (Mathematical Policy Research), and Calvin Luscombe (Amazon), Bruno Martins (Analysis Group), and Xiaoxi Zhao (RAND).

I continue to work with Dr. Kim Geissler, a health services researcher who worked with me as a postdoctoral associate, and we have written a series of papers examining healthcare referral networks (Geissler, Lubin, and Ericson 2016, 2020a, 2021b, 2021) and their association with healthcare organizational structure, insurance contracts, and practice style. This descriptive research has been valuable in establishing facts and inspiring more intensive economics research papers. I am currently advising Dr. Angie Acquatella in a collaboration studying prescription drug payment policy.

## Service

I currently serve as the Faculty Director for Business Analytics in the Undergraduate Program. I led the creation of a new inter-departmental undergraduate business analytics concentration, developed a new curriculum for analytics concentrators, and oversaw the revision of the analytics courses required of all undergraduates in the school. I have shaped the school's strategy toward data and statistics, both in curriculum and in hiring. I have led search committees for cross-departmental hires at Questrom. I served as a member of the Dean's Strategic Council and worked on the team that redesigned our MBA core curriculum. I co-chaired the undergraduate curriculum redesign task force, which examined the future of our undergraduate program and made recommendations to the faculty (under consideration) for changing our curriculum to be both more rigorous and more flexible.

I have also contributed leadership more broadly in the academy. My work as a journal co-editor of both the *Journal of Public Economics* and the *American Journal of Health Economics* shapes the direction of academic research. The recent launch of the latter journal has enabled me to shape its strategy and test new policies, such as a speedier prior review process and an explicit conflict of interest policy. I have organized and chaired sessions at various conferences, including the American Economic Association meetings and ASHEcon.

As part of my commitment to diversity and inclusion, I have consciously chosen a diverse representation of individuals in case studies and examples in teaching. I have also been an active participant in discipline-based groups that facilitate diversity, such as the American Economic Association's Committee on the Status of LGBTQ+ Individuals in the Economics Profession, and I have served as a mentor for junior faculty in the AEA's LGBTQ+ committee.