# Keith M. Marzilli Ericson: Personal Statement

May 2016

## Research

### *I. Introduction*

I have two often intersecting streams of research: *healthcare markets* and the *psychology of decision-making*. As an applied microeconomist, I have shown how choice and regulation interact in new insurance markets, notably Medicare Part D and the ACA's insurance exchanges. As a behavioral economist, I have made contributions to understanding foundational aspects of decision-making. My research on the interaction of individual decision-making, firms, and regulation has implications for public policy—e.g., how should pricing be regulated on insurance exchanges? (Ericson and Starc 2015, *Review of Economics and Statistics*)— as well as for managerial practice—e.g., how should deadlines and reminders be optimally designed (Ericson forthcoming, *Journal of the European Economic Association*).

My research in healthcare markets can be organized into two main themes: consumer choice, and the response of firms to consumer decision making. In consumer choice, I have pioneered the study of how consumers choose on health insurance exchanges: I identified consumers' price sensitivity and willingness-to-pay for enhanced network access, studied the standardization of policies as a choice architecture intervention designed to increase competition, and compared the effectiveness of insurance mandates versus other financially equivalent policies. I have also made important contributions to understanding how insurers respond to consumer decision-making, showing "invest-then-harvest" pricing in Medicare Part D, the role of renewal defaults, and the interaction of modified community rating regulations with imperfect competition.

My research in behavioral economics has made contributions to the study of reference points for loss aversion, intertemporal choice (both time preference and memory), and the role of privacy in truthful reporting. In an influential article (Ericson and Fuster *Quarterly Journal of Economics* 2011), I tested the theory that reference points are determined by expectations about the future, and linked expectations-based reference points to the endowment effect. For intertemporal choice, I showed that heuristic models make better out-of-sample predictions than utility models for money earlier-or-later choices, and collaborated with neuroscientists to examine how the brain discounts real consumption rewards. I provided the first evidence on memory and "forgetting to act" in the experimental economics literature, linking decisions in the lab to field behavior to show that individuals' beliefs about memory were overconfident. I then examined how memory interacts with present-biased time preferences, deriving implications for the empirical measurement of time preference and the optimal design of reminder systems and deadlines. Finally, I examined how a privacy-protecting survey method affected how frequently sensitive answers were reported. This work not only shows that existing work substantially underestimates both anti-gay bias and the size of the lesbian, gay, and bisexual (LGBT) population, but also provides a service to other researchers by validating the method using a large scale (N=8,709) placebo test.

There is evidence that my research has had an impact on both academic work and policy-making. I have not only published in top economic journals, but also in *Psychological Science* (the highest ranked empirical journal in psychology) and *Management Science* (the flagship management journal). As of May 2016, my work has received 1062 citations according to Google Scholar, and 416 citations in the ISI Web of Science. My work has received media coverage in the *New York Times*, *Washington Post*, *The Atlantic*, *Time*, *Slate*, *Los Angeles Times*, *Boston Globe,* and others.

I am recognized as an expert on health insurance markets, particularly health insurance exchanges and the behavioral economics of health insurance. I have been invited to present my work at conferences focused at policy-makers and practitioners (e.g. Federal Trade Commission, Department of Justice, Academy Health, Society of Actuaries), and am a regular attendee and presenter at the invitation-only NBER insurance and health conferences. I have also received healthcare-related grants and been invited to review healthcare-related grant proposals (e.g. by the Robert Wood Johnson Foundation). I have also had a paper proposal (joint with Justin Sydnor) accepted by the *Journal of Economic Perspectives* that will synthesize the implications of behavioral economics for the study of health insurance.

My core contributions to our understanding of intertemporal choice (both time preference and memory) and to reference points have established my reputation as a leading behavioral economist. I have presented at and am a regular participant at the field's two invitation-only core conferences (the Behavioral Economics Annual Meeting, and the Stanford Institute for Theoretical Economics Psychology and Economics conference), and was invited to spend a week in Berlin at the European Behavioral Economics Meeting. I am also an associate editor at the *Journal of the European Economic Association,* where I handle manuscripts related to behavioral economics, as well as those to relevant to my healthcare research.

## *II. The Economics of Healthcare and Health Insurance*

As an applied microeconomist, my research has examined the intersection between individual decision-making, health insurance, and healthcare markets. Healthcare comprises about 17% of U.S. GDP, has large impacts on human welfare through quality and length of life, and presents unique economic and managerial challenges. Recent changes in healthcare policy and practice have expanded the role of consumer choice in insurance, and insurance plans shape the incentives that patients and healthcare providers face. Insurance plans themselves have been the focus of major public policy initiatives and raise many novel economic questions, such as the interaction between selection and imperfect competition.

### *II.A Consumer Choice in Health Insurance Exchanges*

A major contribution of my research has been to examine the health insurance exchanges—also known as marketplaces— which form the core of the 2010 Affordable Care Act. I was the first to examine consumer choice on health insurance exchanges. I did so by obtaining data from the Massachusetts exchange: created in 2006 as a result of state-based health reform, it served as a model for the federal legislation. In a series of papers on this topic, I have investigated how these markets work and should be regulated.

The insurance exchanges are new markets, about which little was known. Understanding what people chose and valued on the exchange, and how that affected competition, was a high priority question. While the exchanges were intended to drive down prices through the discipline of consumer choice, this requires consumers to choose effectively. In Ericson and Starc (2012 *American Economic Review Papers and Proceedings*), I examined how price sensitive individuals were when purchasing insurance on the exchange. The level of price sensitivity is important, because it governs the ability of firms to charge markups over cost. I found substantial heterogeneity in price sensitivity, which highlighted that competition at the top and bottom of the market may be quite different. I also found a large gain to being the cheapest plan, consistent with a group of very price sensitive enrollees who placed very little value on more generous insurance plans or else relied on ad-hoc heuristic rules (e.g. "simply choose the cheapest plan").

In Ericson and Starc (2015 *American Economic Review Papers and Proceedings*), I turned to the question of *narrow network* insurance plans, which only provide coverage at a small number of hospitals and physicians. Narrow networks present a policy conundrum: while they enable lower cost insurance plans, it is difficult to observe the quality of plans' networks. In my research, I studied the value consumers placed on network generosity, bringing together data on exchange enrollment, claims from the Massachusetts All-Payer Claims Database, and extensive hand-collected network information. I was able to estimate the dollar value that consumers placed on enhanced network access, as well as how this varied by demographics. I also showed that access to certain "star hospitals" was valued more highly than access to other hospitals. My results also suggested that while consumers valued differences in networks between different brands, they had difficulty comparing multiple networks offered by the same brand. In follow up research, I am directly examining the degree to which consumers value insurance brands themselves, as distinct from the underlying generosity of their networks.

In Ericson and Starc (2013 "*How Product Standardization Affects Choice*"), I examined a policy that aimed to facilitate consumer choice: standardization of health insurance plans, so that all plans in a "tier" (e.g. bronze, silver, and gold levels of generosity) had the same cost-sharing parameters. Standardization is relevant not only for health insurance, but for many complex contracts, such as cell phone plans and credit cards. While standardization is believed to make it easier for consumers to compare products and thereby increase price competition in the market, empirical evidence on the effect of standardization was lacking.

In this paper, I used a natural experiment on the Massachusetts insurance exchange to assess how standardization affected health plan choice. I showed that as a result of standardization, both more generous plans and very different brands were chosen by consumers. I developed a model to distinguish the two different channels by which standardization can have an effect: the availability effect (the set of products available to consumers changed pre- and post-standardization) and the valuation effect (the weights consumers placed on attributes changed pre- and post-standardization). The increased generosity in plans chosen resulted from the increased weight consumers placed on plans' cost-sharing parameters such as deductibles (the valuation effect), while the brand shift was due to the availability effect. Surprisingly, we found no evidence of increased consumer price sensitivity. These results imply that standardization may not achieve the goal of increasing price competition in insurance markets, but

that standardization is not a neutral policy either: depending on how it is designed, it can affect the type of plans consumers choose.

An individual mandate to purchase insurance was a cornerstone of the 2010 Affordable Care Act, and policy-makers believed that a mandate and its associated fine would increase enrollment in the exchanges in a way that an equivalent tax would not. Yet while there was a large literature examining framing effects in laboratory-based tasks, there was no evidence on whether mandates would actually be more effective than a financially equivalent tax. In Ericson and Kessler (2016 *Journal of Economic Behavior and Organization*), I examined the decision to purchase health insurance under two different, but financially equivalent policies: one policy that taxes individuals who do not purchase insurance, and another policy that mandates that individuals purchase insurance and fines those who do not comply. I found that prior to the controversy over the mandate's legitimacy, individuals in a hypothetical choice experiment were more likely to purchase insurance when the policy was articulated as a mandate with a fine than when it was articulated as a tax. I documented the timing of the public controversy about the mandate (concurrent with the Supreme Court case), and showed that the controversy undermined the legitimacy of the mandate and highlighted its equivalence to a tax. These results demonstrate that while the framing of government policy matters, framing is not entirely controlled by policy-makers. Moreover, these results suggest that enrollment in the insurance exchanges would have been higher without the political attacks against the mandate. This work has been cited in estimating the effect of the ACA's tax penalty.

Measuring consumers' risk aversion is key input into predicting how policies will affect insurance markets and evaluating their welfare consequences. In Ericson, Kircher, Spinnewijn, and Starc (2015 *"Inferring Risk Perceptions and Preferences…"*), I showed how to identify consumers' risk preferences even if they have incorrect beliefs about their healthcare usage. Demand for insurance can be driven by high risk aversion ("risk preferences") or by high anticipated probability of making an insurance claim ("risk perceptions"). The literature typically assumes that consumers have rational expectations about the probability they will make a claim, uses claims data to identify these risk perceptions, and then estimates risk aversion from the choice between insurance plans. 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. I 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 (e.g. for researchers studying health insurance exchanges), as well as used to evaluate the accuracy of risk perceptions when claims data is available.

In a paper currently in progress (Ericson and Sydnor, in progress), I am studying the limits of what we can learn about consumers' preferences from their choices in health insurance markets. Here, I am examining how liquidity constraints can result in a seemingly irrationally high willingness-to-pay to lower deductibles.

## *II.B Responses of Firms to Consumer Decision-Making in Health Insurance Exchanges*

### *Modified Community Rating*

Health insurance markets are heavily regulated, and this regulation is in flux. In Ericson and Starc (2015 *Review of Economics and Statistics*), I examined how pricing regulation interacts with imperfect competition in the exchanges. Many policy-makers' intuitions and projections about how regulation affects market outcomes relies on perfect competition, in which firms price at cost. However, recent research has shown that the insurance market is imperfectly competitive, enabling insurers to charge markups over cost. 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, these regulations also change markups. The key intuition is that firms price to the marginal consumer, and their first-order condition places more weight on the marginal cost of the more price-sensitive consumer.

I showed empirically that younger, healthier, lower-cost individuals are more price sensitive than older, sick, higher-cost individuals. As a result, when modified community rating links the prices of the healthy and sick, the costs of the healthy get more weight in the optimal pricing condition, lowering the average price that the firm charges. These results have direct implications for age-based pricing regulation. Federal regulations for the insurance exchanges limit prices for the oldest to be no more than three times the price for the youngest. Apart from any potential equity issues raised by this regulation, our work shows that the regulations have substantial efficiency benefits (on the order of $300 per enrollee per year) by lowering the total markups that firms charge for all consumers.

### *Responses to Inertia*

Consumers purchase health insurance plans on an annual basis, but their initial decisions can have long-lasting consequences, as individuals face choice frictions that lead to inertia in product choice. While many papers have shown inertia in choice, my work has examined how inertia affects health insurance markets. Ericson (2014 *American Economic Journal: Economic Policy*) was the first paper to examine how insurance firms respond to individuals' inertia. I examined the launch of Medicare Part D, which provides prescription drug insurance for the elderly through an insurance exchange. I first documented inertia in this market, using a regression discontinuity design for default assignment rules to show that being randomly assigned a particular plan in the first year had persistent effects on choices years later.

Rational firms respond to inertia when setting prices, which is typically seen in the form of introductory offers, which allow firms' to charge a higher price to repeat customers who are likely to be inertial. However, market design decisions shape how firms can respond to inertia; in the Medicare Part D market, introductory offers are legally prohibited. I showed that firms attempt to replicate introductory pricing by entering the market with low prices to acquire market share. In later years, incumbent firms then raise prices on consumers who are less price sensitive, while 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.

This work has implications for the Affordable Care Act's insurance exchanges. It not only predicts large price increases in the market's early years, over and above the rate of healthcare cost increases, but it

also establishes that contract restrictions play a major role in determining market equilibrium. If firms were allowed to charge introductory prices for first-time enrollees, they would choose to do so. If firms cannot, they would like to introduce new products targeted at first-time choosers, while raising the prices on existing products. In Medicare Part D, however, firms were restricted in the number of products they could offer, creating incentives for new firms to enter the market in later years.

Renewal default rules are "nudges" that affect inertia. In Ericson (2014 "*Dynamic Defaults…*"), I showed how default rules affect market equilibrium. I examined the rules that determine what happens when consumers enrolled in an insurance plan do not make an active decision. If a consumer who previously purchased the product does not make an active decision, the default is often to automatically renew their purchase of the product; however, in some markets, the default is 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 elasticity of repeat demand, and thus the pricing pattern of firms. 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. Default rules, in effect, have externalities, and the privately optimal default rule will not typically coincide with the socially optimal default rule. I then showed how the socially optimal default can be determined from simple choice experiments. The Federal Trade Commission is particularly interested in this research because of their role as an enforcer of competitive policies and they invited me to present it at their offices.

### *II.C. Healthcare and Organizations*

Healthcare is delivered by teams of providers, including primary care physicians, specialists, and nurses. In Agha, Ericson, Geissler, Lubin and Rebitzer (in progress, "*Coordination in Teams and the Cost of Healthcare*"), I am examining the organization of teams in healthcare delivery. I modeled team formation via primary care physicians' decisions to invest in working relationships with specialists. Relationship-specific investments are costly, but increase team quality, lowering healthcare costs and increasing quality of care. The model shows that coordination is more likely when physicians concentrate their referrals on a small group of specialists. Using both Medicare data and the Massachusetts All-Payer Claims Database, I show empirically that a primary care physician's referral concentration has an economically large and statistically significant association with their patients' costs. This association persists not only when controlling for patient's health status, but also when controlling for the particular specialists seen. The results are identified, for instance, by comparing two patients with the same healthcare conditions seeing the same cardiologist but coming from primary care physicians that vary in how closely they work with that cardiologist. Better team formation also provides a plausible mechanism through which public policy initiatives, such as Accountable Care Organizations and narrow networks, might affect healthcare.

It well known that there is wide price variation among healthcare providers—some hospitals negotiate very high payment rates with insurers, while others are much lower. In ongoing work (Ericson and Starc, in progress), I am examining price variation between insurers at the same provider—some private

insurers pay a higher rate to a given hospital than others. This work both documents a new fact and explores its implications for bargaining models.

## *III. Foundations of Choice: Behavioral Economics*

### *III.A. Reference points*

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. It has been used to explain the endowment effect, which is the widely observed fact that—in contrast to traditional economic theory— ownership seems to increase the value an individual places on goods. Yet a theory of the reference point is needed to determine what counts as a gain or a loss. Prospect Theory left the reference point unspecified, and it was variously taken to be determined by a variety of factors: the status quo, ownership, social norms, or some other anchor.

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 a complementary but 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. I showed that individuals who were more likely to expect to receive an item valued that item more, as measured by their willingness to trade for an alternative item and via their willingness to accept money instead of the item. Expectations-based reference points have novel implications for firms and policy. For instance, firms' pricing and marketing strategies should attempt to instill in consumers an expectation of purchase, which may then increase their willingness to pay.

This paper was influential because it showed how theories of the reference point could be tested experimentally, as well as how expectations connected to the endowment effect. This experiment was redone in Camerer et al.'s "Evaluating replicability of laboratory experiments in economics" (2016 *Science*), who found very similar point estimates to my original work (high expectations led to a 20-30% increase in value, $p=0.055$ with t-tests, $p<0.01$ in regressions.) 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.

### *III.B. Intertemporal Choice: 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? When do I complete a project relative to the deadline? I have a body of research that advances our understanding of how individuals make such intertemporal choices.

Discounted utility is the primary framework economists have used to analyze intertemporal choice. It encompasses both classical exponential discounting, as well as present-biased models used to study self-control. However, my research indicates that discounted utility is not the best explanation for what participants actually do in the most commonly used intertemporal choice laboratory task. In Ericson et al. (2015 *Psychological Science*), I show that heuristic models of decision-making predict better out-of-

sample than discounting models. In Ericson and Noor (2015 *"Delay Functions as the Foundation of Time Preference: Testing for Separable Discounted Utility"*), I directly test and reject assumptions underlying the discounted utility model.

In Ericson et al. (2015 *Psychological Science*), I examined the "money earlier or later" (MEL) task used in hundreds of economics papers and psychology papers. I proposed that choices in MEL tasks reflected heuristics (rules of thumb), rather than underlying preferences for consumption. The paper proposed a simple heuristic model based on psychological principles (multiple attributes and absolute v. relative comparisons) to explain MEL choices. In the model, when faced with a choice between two options, participants compare the absolute and relative differences of both money and of time. They then attach a weight to each of these attributes. Using a sample of 1000 participants, I conducted the first cross-validated, out-of-sample test of intertemporal choice models. Using modern model comparison methods, I showed that the heuristic models predicted choices out-of-sample much better than any of the discounted utility models in use. I conclude that decisions in MEL tasks should be thought of as reflecting simple rules that may not generalize to intertemporal consumption decisions (such as how much money to spend today). As a result, better methods are needed to examine intertemporal choice.

This work suggests that it will be more productive to measure choice over actual consumption instead of money. While hundreds of studies use the MEL task, my work in McClure et al. (2007 *Journal of Neuroscience*) is one of the few early studies to measure time preference over actual consumption. In this study, participants were deprived of liquids for a few hours before they came to the lab. Participants were then given the choice between drinking squirts of water at various points in time: e.g. 1 squirt now, or 2 squirts in five minutes? Measuring preferences for consumption at the minute-by-minute level, I found evidence of present-bias in time preference.

Regardless of whether actual consumption or money is considered, the time preference literature typically assumes *separable* discounted utility (SDU), in which the discount rate does not depend on the level of money or utility. Even under this restrictive assumption, measuring time preference is challenging. It requires estimating both the discount function for utility and the curvature of the utility function over consumption. However, it is difficult to measure the curvature of the utility function, particularly if researchers allow for non-standard risk preferences or use tasks that involve actual consumption.

In Ericson and Noor (2015 *"Delay Functions …"*), I proposed a new "delay function" method that enables us to directly test the SDU assumption and estimate the shape of the discount function (i.e. is discounting constant/exponential, or present-biased?) without measuring the curvature of the utility function. The typical method asks, in effect, "If you were choosing between $10 today and $*x* in 1 month, how large would *x* have to be in order to make you indifferent?" The delay function method asks, "If you were choosing between $10 today and $15 in *y* months, what would *y* have to be to make you indifferent?" Delay functions thus fix the dollar amounts in the options individuals face and vary the time dimension. As a result, it is not necessary to estimate the utility function for money to determine whether discounting is SDU and whether discounting is exponential or present-biased.

I estimated delay functions on a representative sample of the U.S. population. If, like previous literature, we were to simply assume that SDU held, I would have classified more than half our analysis sample as

exponential discounters. However, once I tested whether the SDU assumption held, I rejected SDU for a majority of the analysis sample in favor of a model in which discounting depends on the level of reward available. Like Ericson et al. (2015 *Psychological Science*), this paper suggests that the standard model of time preference—even with present-bias—does not describe what many participants are doing in experiments. This work leaves open the question of the extent to which the SDU model describes time preference in field choices.

Intertemporal choice is an interdisciplinary field, and my work here has often been in collaboration with neuroscientists and psychologists. Neuroscience can enhance economists' understanding of the underlying decision-making processes. Many recent economics models of time preference posit conflicting systems or selves (e.g. a "planner" versus a "doer," or a patient system versus a system concerned with immediate gratification). In McClure et al. (2007 *Journal of Neuroscience*) we conducted fMRI brain scans of participants while they made consumption decisions. These scans showed that time preference was constructed from the interaction of more patient areas of the brain (which value immediate and delayed rewards roughly the same) and impatient areas (which show much greater activation when an immediate reward is present than when a delayed one is).

Based on this body of work, I wrote a review of the time preference literature for the *Journal of Economic Literature* (Cohen et al. 2015, revise and resubmit). I take the perspective that time preference is fundamentally about when individuals consume utility, which is distinct from how consumption is financed. (For instance, a vacation today could be paid for out of a bank account either today or in three months with some interest; changing the timing of payment is not the same as changing the timing of consumption.) However, the workhorse tool for economists and psychologists studying intertemporal choice has not been a consumption task, but a financing task: the "money earlier or later" (MEL) task. It is a puzzle that MEL tasks have been used to study intertemporal choice, as economic theory predicts that MEL should only elicit time preference under a narrow range of conditions. The review explores the connection between MEL tasks and other intertemporal choices, shows what can be inferred about preferences from MEL under different assumptions, and highlights the differences in findings between MEL tasks and tasks involving real consumption experiences.

### *III.C. Intertemporal Choice: Limited Memory*

Deciding when to complete a task depends not only on time preference, but also on memory: do you remember to act when you plan to act? In Ericson (2011 *Journal of the European Economic Association*), I provided some of the earliest evidence on memory for action in the economics literature. I examined whether people remembered to claim delayed payments, and compared their claim rates to their incentivized forecasts of how likely they were to claim. Despite being able to use any technology (calendars, schedulers, etc.) they desired, only 50% of people claimed a delayed reward. In contrast, individuals' forecasts implied that they should have claimed 75% of the time. People not only forgot, but they were overconfident about the probability of remembering. These results had immediate implications for the design of offers such as rebates—consumers will think rebates are more valuable than they truly are, because they will forget to claim them.

Following my work in Ericson (2011 *Journal of the European Economic Association*), a number of other researchers began to test whether reminders (e.g. text messages to increase household savings rates)

changed behavior. Providing reminders, it turns out, seems to be effective in helping individuals achieve their own goals, including saving and taking medication. This result is puzzling. Why is it effective for researchers to provide reminders, when individuals can easily and cheaply set up reminders themselves?

In Ericson (forthcoming *Journal of the European Economic Association*), I moved beyond studying memory in isolation, and examined how memory interacts with present-biased time preferences. The interaction of these two biases yielded novel results. First, counterintuitively, present-biased individuals can actually benefit from being forgetful: knowing you will forget in the future can serve as a commitment device to act today. Second, it explains why providing unexpected reminders to people can be effective even though cheap reminders were already available: not only do individuals with present-bias procrastinate in setting up reminders, but sophisticated individuals recognize that setting up their own reminders enables procrastination. A key result in the paper shows how to optimally design reminder systems and deadlines, and how that design should vary by whether individuals are present-biased or not. Finally, it shows that empirical estimates of present-bias from task completion behavior will be incorrect if limited memory is not accounted for.

### *III.D. Privacy and Measuring Sensitive Topics: Sexual Orientation*

Having accurate measures of sexual orientation and the extent of anti-LGBT sentiment is important for both research and public policy. Many areas of research use data about the LGBT population, including work on discrimination in the labor market, household labor supply, economics of the family, and public health interventions to reduce sexually-transmitted diseases. However, answers to survey questions about sensitive topics might be biased towards social norms ("social desirability bias"), with respondents giving socially approved answers rather than honest answers.

I tested whether social desirability biased affects existing estimates of the size of the LGBT population and anti-LGBT sentiment in Coffman, Coffman, and Ericson (forthcoming, *Management Science*). 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 makes a contribution to the literature on sexual orientation, but also for measuring sensitive behaviors more broadly: I evaluated the validity of the methodology using the first large scale (N=8,709) placebo test for the Item Count Technique ever conducted. In an experiment with 8 "placebo" statements that should not be affected by social desirability bias, seven of the eight placebo items produce treatment effects that are statistically indistinguishable from zero. This paper is already widely cited and has received substantial media attention (e.g. *The Atlantic, Boston Globe*, and *New York Times*).

# Teaching Statement

I have taught both undergraduates and MBA students, and advised several PhD students. I have taught the required econometrics course for undergraduate business students at Boston University. I redeveloped the course to focus on "making decisions with data," examining how statistical methods can extract insights from data to inform business decisions. Each concept in the course is now linked to a case or business problem. In collaboration with colleagues, I wrote a text to be used in the course, *Making Decisions with Data*.

I have also developed a series of classes on behavioral economics for different audiences. First, I designed an undergraduate honors seminar on the "Art and Science of Decision Making," in which I used active learning, in-class experiments, and in-depth projects to explore key concepts in decision-making. Second, I have taught an MBA elective on "Improving Your Decisions," applying behavioral economics and game theory to personal and business decision situations. Finally, I developed an interdisciplinary course for undergraduates on "The Psychology of Decision-Making: Implications for Business and Public Policy," in collaboration with Peter Blake from the Department of Psychology. Because this course enrolls both psychology, business, and economics students, we developed a series of new experimental materials and new cases. My development of this course received a grant from the Provost's program to encourage interdisciplinary teaching collaborations.

My teaching philosophy is informed by psychological research and relies on active learning. The form that active learning will take will vary by context, but I have strategies to engage students in both large lectures and small seminars. 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 synthesize new ideas 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 (electronically) in advance or being ready to answer questions posed to them during class.

I have also been sought out as a mentor for PhD students from the Boston University Department of Economics. I have advised four students from the Boston University economics department: Julie Shi (now at Peking University), Timothy Layton (Harvard University, Dept. of Healthcare Policy), Emily Gee (U.S. Dept. of Health and Human Services), Sara Marcado (post-doc, London School of Economics/Imperial College).

## Service Statement

I currently serve the profession as an Associate Editor at the *Journal of the European Economic Association (JEEA)*, a highly ranked general-interest journal. I also regularly review papers for top journals economic and management journals, including the *American Economic Review, Quarterly Journal of Economics (QJE)* and *Management Science*, as well as for leading healthcare journals (e.g. *Journal of Health Economics* and *Health Affairs)*. I received an "excellence in refereeing award" from both the *QJE* and *JEEA.* I have also evaluated funding proposals, serving as a grant reviewer for both the National Science Foundation and the Robert Wood Johnson Foundation

I have contributed to organizing conferences and seminars as well. I have organized sessions for the American Economic Association Annual Meeting and the National Tax Association. I have served as a Consumer Choice and Behavioral Economics theme reviewer for the Academy Health Annual Meeting, and on the scientific committee of the American Society of Health Economics Conference. I co-organize both the BU-Harvard-MIT Health Seminar and the Markets, Public Policy, and Law seminar at Boston University.

I also aim to translate insights from research into policy and practice. I am one of the organizers for the Behavioral Science and Policy Association's "Behavioral Economics and Health Working Group." As part of this group, I am producing a proposal for the White House Social and Behavioral Sciences Team about how to apply insights from behavioral economics to healthcare. I have also given talks at the Federal Trade Commission and Department of Justice to inform policy toward healthcare markets.