

## Research Statement

Christopher J. Cronin

My research focuses on health as a form of human capital. People invest in this capital via formal medical care as well as through their behavior. My primary goal is to understand the extent to which policy can promote human flourishing by influencing these health investments, which requires an understanding of both the value associated with health improvements and the incentives that can alter health behaviors.

I have taken two broad approaches in my work. The first approach is prospective, using data to inform models of economic behavior with the goal of testing whether hypothetical policies will alter health behaviors in a way that improves welfare. The second is reflective, considering policies that have been enacted and asking whether particular objectives were achieved, and if not, why? Separated by research approach, I summarize below how my current and future research plans help to achieve the goal described above.

*Prospective approach* – One strand of my research models the incentives driving health behaviors in an effort to test the welfare implications of prospective policies. This work uses dynamic-stochastic decision-making (DSDM) models of consumer choice, a technique used extensively in health and labor economics, as well as in industrial organization, when critical costs or benefits associated with decision making are realized in the future. This characteristic describes virtually *all* health behaviors. The method of inquiry typically involves: (i) specifying a theoretical model of how behavior should unfold over time, i.e., dynamic behavior; (ii) solving each individual's optimization problem as defined by the model; and (iii) using data to estimate unknown parameters of the model. Upon estimating these parameters, the model can then be used to simulate consumer behavior under counterfactual policies; subsequently, welfare and other policy variables can be evaluated.

This method is used in my paper “Insurance-Induced Moral Hazard: A Dynamic Model of Within-Year Decision Making Under Uncertainty,” published in 2019 in the *International Economic Review*. The paper is concerned with insurance-induced (ex-post) moral hazard – health insurance lowers the price of medical care below the market rate, naturally encouraging greater consumption and, therefore, higher premiums. In the paper, I measure the scale of this moral hazard in the United States. The paper features a model of within-year medical care decision making (WYDM), where agents first select a health insurance plan as a function of expected future spending. Agents then make repeated medical care decisions over the course of the insurance year in response to health shocks, weighing the cost of medical care against the productive benefits of treatment. The innovation of the paper is modeling medical care decisions, health transitions, and price evolution as dynamic processes that involve uncertainty. Estimating patient price sensitivity is fundamental in measuring moral hazard. The advantage of the WYDM model is that it captures changes in healthcare choices as expected prices change over the health insurance year as agents approach and ultimately cross their insurance deductible. This method contrasts with a simpler approach commonly taken in the literature, which is to assume individuals receive a single health shock after selecting their insurance plan, followed by an annual medical expenditure (AME) decision.

Though the main contribution of the paper is to design and estimate a more realistic model of consumer behavior, the results are policy relevant. For example, at the time of writing, there was an

important debate in the US regarding the potential costs of the Affordable Care Act (ACA). In the paper, I use the model to predict how much US healthcare expenditures would increase should the 41 million uninsured Americans be able to obtain health insurance because of the ACA. The figure is \$111 billion (about a 3.7 percent increase over observed 2014 spending). Using the simpler AME model, in contrast, the figure is about 10 percent smaller.

A challenge with the paper above is that the focus on health insurance requires that I model *all* healthcare choices, making it difficult to consider costs beyond average prices, or benefits beyond average health improvements for broad classes of treatment, e.g., prescription drugs. In “What Good are Treatment Effects without Treatment? Mental Health and the Reluctance to use Talk Therapy,” (conditionally accepted, *Review of Economic Studies*), Matt Forsstrom, Nick Papageorge, and I use a similar style of DSDM model to analyze a narrow class of treatment choices; namely, the decision to use antidepressants or psychotherapy for the treatment of depression and anxiety.

The motivation for the paper is that the most effective medical treatments often come with the most serious drawbacks, e.g., high prices, significant time costs, side-effects, etc. These drawbacks can prevent one medical treatment from being strictly “better” than another and treatment decisions likely reflect such tradeoffs. In the paper, we summarize the clinical literature measuring the efficacy of these mental health treatments, showing that, for most, psychotherapy is much more effective than antidepressants on average. In spite of the apparent superiority of psychotherapy, individuals are four times more likely to opt for a prescription. We thus set out to understand which costs associated with psychotherapy prevent its use. Because the treatments are narrowly defined, the model can account for the myriad costs associated with the treatments, e.g., out-of-pocket prices, time and employment-related costs, uncertainty related to treatment effectiveness, the possibility that treatment is harmful, etc. We ultimately find that substantially reducing these therapy costs, including the most obvious impediments to use like monetary and time cost, do little to increase use. The most important barrier to use is simply the disutility of therapy, which encapsulates a number of costs that are not measured directly. Examples of these costs include stigma, biases in beliefs, and that therapy actually requires one to “do the work” outside of sessions as well. Understanding the scale of these costs is important and policy relevant and, thus, motivates future work.

*Reflective approach* – I have published several papers that evaluate the impact of existing policies on related health behaviors and outcomes. A unifying theme of this research is my attempt to look at expected as well as unanticipated responses to policies. My research highlights how in many situations, policies alter the incentives for productive health behaviors in unexpected ways, resulting in welfare loss.

An example is the 2004 black box warning label placed on antidepressant packaging in the US. The label warns of an *increased* risk of suicidality associated with taking the drugs for children and adolescents. The research underpinning this warning has always been controversial, made even more so by recent research showing the opposite relationship. In 2007, the original suicidality warning was expanded to include young adults under the age of 24. In “The Labor Market Effects of Antidepressants: Evidence from the 2007 FDA Black Box Warning,” published in 2020 in the *Journal of Health Economics*, Aline Bütikofer, Meghan Skira, and I study the behavioral response to the expanded FDA warning. Our empirical analysis compares individuals with and without a history of depression, before and after the warning was expanded. Examining men and women of different age groups

separately, we find a reduction in antidepressant use for women aged 35-49, a group *not* explicitly targeted by the new warning, but a result consistent with other research showing that targeted informational campaigns can alter the behavior of non-targeted groups. We also demonstrate that the campaign generated a reduction in employment for this group, producing a reduction in labor earnings of \$11.8 billion.

In another example, my co-author Bill Evans and I study the COVID-19 pandemic and policies designed to encourage social distancing. In “Total Shutdowns, Targeted Restrictions, or Individual Responsibility: How to Promote Social Distancing in the COVID-19 Era,” published in 2021 in the *Journal of Health Economics*, Evans and I show that controversial stay-at-home orders passed in the US in the spring of 2020 explain almost none of the large (observed) reduction in foot traffic to higher-risk establishments like restaurants and bars, but explain a moderate share of the reduction in foot traffic to lower-risk establishments like grocery stores and golf courses. The implication is that rational agents avoided risky establishments in the absence of policy, yet policy caused them to further avoid low-risk establishments, possibly reducing welfare. For example, stay-at-home orders reduced foot traffic to outdoor venues, which were some of the safest locations with the lowest transmission rates. We go on to show that in locations without stay-at-home orders, targeted mobility restrictions, like restaurant dine-in bans, were very effective at reducing mobility to the targeted industry without spilling over into others. These two sets of results indicate more targeted policies would have been much more effective.

The most serious consequence of COVID-19 pandemic has been the substantial death toll, which the World Health Organization estimates at almost seven million world-wide to date. Some of my research focuses on how both personal and policy responses to COVID-19 led to an excess of non-COVID deaths (i.e., deaths not directly caused by a Coronavirus infection) in 2020 and 2021. In “Nursing Home Quality, COVID-19 Deaths, and Excess Mortality,” published in 2022 in the *Journal of Health Economics*, Evans and I document that while the highest quality US nursing homes witnessed fewer COVID-19 deaths than the lowest quality nursing homes, they had substantially more non-COVID and even total deaths. We further document that the highest quality homes were the most likely to follow CDC guidelines on every measurable dimension: more safety equipment, more testing, higher vaccination rates, fewer staff shortages. As such, we argue these homes also followed CDC guidelines on isolating patients from outside visitors and one another, producing loneliness, despair, and, in the worst cases, the non-COVID deaths that we document. This research was supported by NIH grant #R21AI164391.

In a related paper, “Excess Mortality from COVID and non-COVID Causes in Minority Populations,” published in 2021 in the *Proceedings of the National Academy of Sciences*, Evans and I measure total US non-COVID excess mortality (i.e., pandemic era non-COVID deaths in excess of the historical average) in 2020 at nearly 51,000. We were among the first to measure these deaths, and to highlight that non-COVID excess mortality was not spread equally across gender, race, and age groups. We show that while minorities represent 36 percent of 2020 COVID-19 deaths, they represent 70 percent of non-COVID excess deaths and 58 percent of non-COVID excess life years lost. Outcomes are worst for Black, non-Hispanic men, for whom nearly half of the excess life years lost in 2020 were due to non-COVID causes.

*Ongoing work* – I have several ongoing projects. First, an interesting empirical finding from my mental health treatment project is that many individuals attempting therapy in our data drop out after just a

few sessions. This practice of “discontinuation” is well documented in the psychology literature, but is poorly understood. Ethan Lieber, Meghan Skira, and I are currently in the pilot phase of a randomized controlled trial (RCT) that pays military veterans entering therapy for reaching several milestones; namely, completing 6, 12, and 18 sessions. We plan to use the data to estimate the impact of the program on therapy continuation, as well as the impact of therapy intensity on outcomes, such as employment, earnings, and long-term mental health treatment use. Moreover, we plan to use the data to estimate a DSDM model, using the exogenous variation in the returns to treatment created by the RCT in identification. This project is done in collaboration with the Wilson Sheehan Lab for Economic Opportunities (LEO) where I am a faculty affiliate. Lieber and I have other work supported by LEO, including “Employee Demand for Skills Training: Experimental Evidence from Home Caregivers,” (revise and resubmit, *Journal of Health Economics*). These collaborations represent investments in my own human capital, in effect “training” me as a designer and conductor of RCTs, which will open future research opportunities.

Another set of ongoing projects leverages a unique data set containing eight years of leave decisions by nearly 1,000 US public school teachers. I created this data set from several administrative sources. The data contain two novel features: (i) a daily record of leave behavior by type (i.e., sick, personal, emergency, or unpaid leave) and (ii) the daily running balance of available leave. To my knowledge, no other economist has used a dataset of American employees with these two features. In “The Anatomy of U.S. Sick Leave Schemes: Evidence from Public School Teachers,” (under review, *Review of Economics and Statistics*), Matthew Harris, Nicolas Ziebarth, and I use the unique features of these data to, among other things, estimate how changes in a teacher’s leave balance affect their leave use. We are the first to estimate this relationship, which is critical for understanding the optimal design of sick-leave schemes. We estimate that a 10-percent increase in leave balance increases leave use by 4.5 percent on any given day; however, we show this relationship is strongest at the bottom of the balance distribution, implying that teachers stop taking leave as they approach a balance of zero to avoid taking unpaid leave. We go on to provide evidence that low balances lead to presenteeism (i.e., working sick) and document further evidence of spillovers (i.e., low balance-induced presenteeism leads to illness spread within the school).

The same co-authors and I intend to use the data described above to study time preferences, a topic of fundamental importance in modeling any dynamic healthcare choice. For example, one reason 11 percent of the adult population still chooses to smoke may be that they simply do not value their future health relative to their current enjoyment of cigarettes. It is well known in the DSDM literature that estimating from data the degree to which agents “discount” the future is challenging. As a result, in most papers, including my own, the parameter that controls the level of discounting is simply selected by the researcher based on professional norms, rather than estimated from data. In the context of our teacher data, accumulated leave carries significant monetary value at the time of retirement, and retirement eligibility is largely determined by experience. Thus, differences in leave use by teachers of slightly different experience levels, but who are otherwise observationally equivalent, can shed light on how much teachers value these future payments. We plan to use this idea to elicit a measure of discounting within a DSDM model of sick leave and retirement decisions. Once estimated, we can use the model to predict teacher responsiveness to alternative sick leave and retirement schemes that might be equally expensive but welfare improving.

*Research Impact* – Finally, media coverage as well as seminar and conference invitations offer evidence of my research impact. My work has been covered by national news outlines like the New York Times,

National Public Radio, Bloomberg Businessweek, and Marginal Revolution. I have been invited to present my research at numerous department seminars, including at the University of Pennsylvania, Johns Hopkins University, Washington University in St. Louis, Rice University, the University of North Carolina at Chapel Hill, and others. My research has been selected for presentation at many of the best and most selective economics conferences, including the Dynamic Structural Econometrics Conference, both American and European Health Econometrics Workshops, the Summer Forum at the Barcelona School of Economics, and multiple NBER conferences.