

RESEARCH STATEMENT
Developing Insights to Improve Health Policies
By Targeting Health Spending to Vulnerable Groups

Amanda Kowalski

May 8, 2023

I am a health economist. My most influential contributions thus far use novel approaches to evaluate the impacts of health policies on vulnerable groups. By examining long-term economic outcomes in administrative tax data, my research finds that subsidized health insurance for [low-income children](#) partly pays for itself. By incorporating key features of the Affordable Care Act and the Massachusetts reform that preceded it into canonical theory and taking it to the data, my research finds that an individual mandate makes health insurance more affordable for [the uninsured](#) as well as the insured. By comparing discontinuities in spending and mortality around a clinical cutoff, my research finds that spending on [at-risk newborns](#) delivers high returns.

Beyond describing policy impacts, my research examines tradeoffs at policy-relevant margins to understand overall impacts on society and to inform the design of optimal policies. By examining the relationship between mammography behavior and long-term rates of breast cancer in a large clinical trial influential to mammography guidelines, my recent research finds that [women more likely to receive mammograms](#) are healthier and more likely to be overdiagnosed with breast cancers that would not cause symptoms for decades. Applied to current guidelines in the United States, my findings imply that an improvement in the targeting of mammograms would provide a rare opportunity to decrease health spending and improve health.

In 2019, I was awarded the [ASHEcon medal](#) for “an economist aged 40 or under who has made the most significant contributions to the field of health economics.” I have also received the [NSF CAREER Award](#) and the [Yale Greer Prize](#) for my body of work. I have 4 major publications, the most recent of which is [sole-authored](#), in top general interest journals in economics: the [American Economic Review](#), the [Review of Economic Studies](#) (2), and the [Quarterly Journal of Economics](#). Of my 16 published papers and 5 papers at earlier stages, only one does not explicitly discuss a topic within health economics. That paper introduces an estimator that I co-developed for a separate application to health economics.

A deep interest in health policy has been integral to my contributions because health policy relies on a complicated patchwork of institutions. To make my findings on mammograms useful to current guidelines, I examined how previous results from clinical trials informed those guidelines with guidance from participants in the literature. I shared what I learned separately in the [Journal of Economic Perspectives](#). Most importantly, long-term results from large clinical trials do not provide compelling evidence of reductions in all-cause mortality, but they do provide compelling evidence of overdiagnosis, which motivates my focus on overdiagnosis in my main work.

One theme in my research process is that I make connections between useful advances in other fields and long-standing questions in health economics. I aim to make contributions that represent more than the next incremental step by using these advances. I began the process while studying public economics and econometrics as a graduate student interested in health. Since then, I have been deliberate in seeking connections. I have done quite a bit of academic travel, and I have held several visiting positions that have been instrumental in helping me to make connections. While I was visiting the Brookings Institution in Washington DC from 2011-2012, shortly after learning about cutting-edge research in public economics that used administrative tax [data that were](#)

[notoriously difficult to access](#), I started a collaboration with David Brown and Ithai Lurie at the US Treasury. Their access to the administrative tax data was crucial to our ability to extend [seminal work](#) from health economics and find that [expansions of Medicaid to low-income children decreased mortality, increased college attendance, decreased fertility in early adulthood, and partly paid for themselves through increased taxes and decreased transfers by age 28](#).

Beyond connections to data, connections to models from public economics have been crucial to my ability to characterize tradeoffs, estimate the impacts of counterfactual policies, and inform optimal policies by estimating impacts on societal welfare. In early work, I adapted a [canonical nonlinear budget set model of taxation](#) to [quantify a key tradeoff](#): employer-sponsored health insurance offers protection from risk but also induces consumption of extra care. In later work, Jonathan Kolstad and I adapted a [canonical model of mandated benefits](#) to incorporate the elements of the Massachusetts and national reforms most important to the labor market. The model allowed us to [extrapolate](#) from the Massachusetts reform and use graphical deadweight loss triangles to estimate the size of the distortion to the labor market. In a different paper with Martin Hackmann, we extended an [influential model of adverse selection](#). We found that [the individual mandate under the Massachusetts reform improved welfare in the individual health insurance market by decreasing the distortion from adverse selection, but the larger penalty under the national reform was closer to optimal](#).

Connections to econometrics have allowed me to answer long-standing questions from health economics in new ways and to pursue new areas of inquiry within health economics. Douglas Almond, Joseph Doyle, Heidi Williams, and I were able to find that [marginal returns to medical care on at-risk newborns are high relative to conventional benchmarks](#) by recognizing that the widespread use of cutoffs in clinical guidelines enables the use of regression discontinuity designs. In my recent work, I combine advances from the econometric literature on treatment effects with the model that I used to examine [adverse selection](#) in insurance markets to model selection and treatment effect heterogeneity within experiments. The model allows me to [reconcile findings](#) from the Massachusetts health reform and the Oregon health insurance experiment by showing that the Oregon experiment expanded coverage to sicker people who were more likely to increase their emergency room utilization upon gaining coverage. It also allows me to use data from an influential clinical trial to [ask whether current guidelines target mammograms appropriately and find that they do not](#).

My work with the greatest impact has two features. First, it has a tight connection between an important question and an empirical strategy. Second, it presents ideas simply. In my current work, I have increased my ability to forge tight connections between questions and empirical strategies by collaborating with others whose expertise complements mine in terms of subject matter and methodology. I also continue to invest in making complicated ideas simpler by presenting them [graphically](#). Simplicity exposes my assumptions to more scrutiny such that the assumptions that survive the process are more compelling. In turn, my work is more useful in health economics and more accessible to researchers in other fields, clinicians, and policymakers.

By teaching others about my research process, I aim to have a broader impact. I help students to develop a “research mindset” through participatory exercises. With undergraduate students, I have developed a series of [problem sets](#) that teach students how to replicate and think critically about research based on what I have learned in my own work. I circulate the problem sets on my website, and I have shared the answer keys with faculty at many institutions. Outside of formal classrooms, I have led a [research team](#) that has included full-time research assistants for the past ten years. I

help the members of my team to grow as researchers by sharing research projects holistically, not just through isolated tasks.

Here, I aim to do more than summarize the findings of my individual papers. I discuss what motivated them, what their key innovations were, how those innovations enabled their findings, and how they have had an impact. I begin by discussing three selected major publications. I then discuss my portfolio of related work, providing additional detail on how my papers build on each other in terms of substance and methodology. I conclude by discussing my recent and current work, including a current project on the equity consequences of targeting Covid-19 vaccines through a series of randomized lotteries.

SELECTED MAJOR PUBLICATIONS

Health Insurance: Long-Term Impacts on Low-Income Children

The greatest legacy of the Affordable Care Act will likely be through the long-term impact of state-level expansions of subsidized health insurance to low-income households through Medicaid. The long-term impact of previous expansions to Medicaid can inform the potential long-term impact of current expansions. With coauthors from the US Treasury, I find that childhood Medicaid expansions for children born from 1981 to 1984 partly paid for themselves in the long term through higher tax revenues and lower tax credits (Brown, Kowalski, and Lurie, 2020, [“Long-Term Impacts of Childhood Medicaid Expansions on Outcomes in Adulthood,”](#) *Review of Economic Studies*).

The main advance that facilitates our findings is access to the population of administrative tax data, which allow us to examine impacts on important long-term outcomes. One challenge in using administrative tax data for research in health economics is that it contains minimal information on health insurance. We construct childhood Medicaid eligibility in the data using linked information from parental tax forms and a [Medicaid calculator](#) that we compiled from many sources, including historical documents on state policies. Effects of Medicaid can be difficult to detect because Medicaid targets the poor, who could have worse outcomes despite improvements through Medicaid. Following [seminal work](#) in health economics, we address this challenge using a “simulated instrument” strategy that isolates policy variation across states and birth month cohorts. We build confidence in the strategy using a dose-response exercise that uses longitudinal variation in childhood income.

We present our results in figures that show the age profile of impacts from age 19 to 28. The figures help us to illustrate mechanisms behind impacts on taxes. Medicaid eligibility increases college attendance and decreases fertility in early adulthood, followed by improvements in economic outcomes. Decreases in mortality accumulate over time. Given these benefits, the main tradeoff in expanding Medicaid is the cost. Rather than obtaining a cost estimate from the literature, we estimate impacts on costs using the same empirical strategy that we use to estimate impacts on benefits with historical data on Medicaid spending. Discounting total taxes and Medicaid spending at a 3% rate, we divide our estimated increase in total taxes by our estimated increase in Medicaid spending to find that the federal government saves 58 cents for each dollar that it spends on childhood Medicaid by the time those children reach age 28.

Putting our findings in the context of findings on a wide variety of government policies, subsequent [high-profile research in public economics](#) concludes that along with investments in education,

investments in Medicaid for children have historically delivered the largest marginal value of public funds. As states have considered whether to implement Medicaid expansions under the Affordable Care Act, the popular press has covered our [findings in several outlets including the New York Times](#). It has also influenced projections by the [Congressional Budget Office](#). Our work was recognized in 2021 as a [finalist for the NIHCM research award](#).

Health Insurance: Welfare Impacts of Insuring the Uninsured in an Adversely Selected Market

One of the most controversial aspects of the Affordable Care Act (ACA) was the individual mandate, which required individuals to have health insurance or pay a tax penalty. The penalty was not implemented until 2014, but the state of Massachusetts implemented a similar penalty in 2006. At the time, Massachusetts was one of a handful of states that already had regulations established by the ACA that limited the ability of health insurers to deny coverage and set prices based on health risk. Such regulations could induce adverse selection of the sickest people into the insurance pool. [Canonical theory](#) held that a mandate could address the problem of adverse selection by drawing healthier people into the pool. The most exciting aspect of the theory was that the mandate had the potential to make insurance more affordable to everyone; it could alleviate the tradeoff between making health insurance more affordable for the healthy or the sick.

Jonathan Kolstad and I began a series of projects on the 2006 Massachusetts health reform just after its passage when we were graduate students in Massachusetts. Our most important contribution is our work on adverse selection, coauthored with my former student Martin Hackmann (Hackmann, Kolstad, and Kowalski, 2015, [“Adverse Selection and an Individual Mandate: When Theory Meets Practice,” American Economic Review](#)). [Influential previous work](#) on adverse selection had found a negligible welfare impact of adverse selection in a context in which it affected the affordability of more generous employer-sponsored coverage. We find a meaningfully large welfare impact in a context in which it affected the affordability of *any* coverage. We also recover the optimal individual mandate penalty, which is higher than the penalty in Massachusetts and closer to the initial penalty established by the ACA.

The key innovation that enables our findings is our linkage between important elements of the Massachusetts reform and a transparent graphical model of adverse selection. [Previous work](#) had estimated welfare in insurance markets using variation in prices. We use variation induced by the Massachusetts reform through the establishment of the individual mandate and the health insurance exchange. Although changes in prices induce a walk along the demand curve, the individual mandate induced a shift in the demand curve by the amount of the penalty, which allows us to identify its slope. The enrollment of healthier people into insurance coverage decreased the average costs that insurers paid on their behalf, which allows us to identify adverse selection. In our innovative data on the individual health insurance market, we can observe average premiums as well as average costs, which allow us to identify a decrease in markups induced by the establishment of the exchange and its associated welfare impact.

To isolate changes induced by the reform from changes that also occurred in other states, we estimate difference-in-difference regressions for enrollment, costs, and premiums with the synthetic control method. We feed the estimated coefficients directly into our model as sufficient statistics, which allows us to construct an empirical analog of our theoretical graph. We recover the welfare gain from the reform as the area of a graphical region. In this region, the marginal cost of insuring the uninsured is less than their willingness to pay for coverage, but they went uninsured before the reform because premiums based on the average cost of sicker enrollees were higher than

their willingness to pay. By drawing healthier enrollees into the pool, the individual mandate decreased premiums for everyone while increasing coverage and thereby increased overall welfare. Our paper was awarded the [NIHCM Research Award](#) and was highlighted in the announcement of my [Yale Greer Prize](#). A student who completed the problem set on this paper conducted [successful dissertation research on adverse selection in the market for natural disaster insurance](#).

Health Spending: Marginal Returns to Spending on At-Risk Newborns

Health spending now represents almost a fifth of all spending in the United States, making the tradeoff between increased spending and increased health an important one to study, especially on policy-relevant margins. However, the marginal return to health spending is difficult to study because patients in worse health often receive more care. Douglas Almond, Joseph Doyle, Heidi Williams, and I propose a novel approach to estimate the marginal return to health spending and apply it to spending on at-risk newborns (Almond, Doyle, Kowalski, and Williams, 2010 [“Estimating Marginal Returns to Medical Care: Evidence from At-risk Newborns,”](#) and [2011 comment, *Quarterly Journal of Economics*](#)).

Low birth weight is salient to clinicians, and its costs have received attention in [high-profile research](#) in economics. Our innovation is to recognize that low birth weight is not just a health outcome: it is an input into decisions about health spending. More broadly, the wide use of diagnostic thresholds by clinicians creates discontinuities in spending for people in similar health at policy-relevant margins. Thresholds that affect spending on at-risk newborns are important, especially because of increasing use of high-cost technologies. We compare newborns just above and below the “very low birth weight” threshold of 1500 grams (just under 3 pounds, 5 ounces).

We present our main results in two simple figures that demonstrate visible discontinuities without the need for superimposed trend lines. First, we plot mortality by birth weight using the census of available United States birth certificate data from 1983 to 2002. Even though mortality generally decreases as birth weight increases, newborns just below the threshold are less likely to die. Second, we plot spending and length of stay by birth weight using data from all hospitals in several states. Newborns just below the threshold have higher spending and longer hospital stays. Because we do not have data on birth weight, spending, and mortality in the same dataset for all newborns, we use a two-sample instrumental variable strategy to obtain our main result. We divide the estimated discontinuity in spending from one sample by the estimated discontinuity in mortality from another to arrive at the marginal cost of saving a newborn life around the threshold. Our estimate is small relative to conventional benchmarks, demonstrating high marginal returns at the cutoff. Our results imply that the optimal cutoff would be higher.

Building on our main analysis, we show that discontinuities are more pronounced in low quality hospitals, which appear more responsive to cutoffs. Moreover, there is a dose-response relationship between the impacts on spending and mortality across hospitals with different levels of quality. This relationship builds confidence that our main results identify a causal relationship between spending and mortality. Overall, our findings build confidence in [previous high-profile estimates](#) of the high returns to medical spending that rely mainly on variation over time.

Our work has received attention from clinicians through the [Garfield Economic Impact Award](#). It also received the [HCUP Outstanding Article of the Year Award](#). It has inspired many subsequent

studies, including a [high-profile extension that finds returns to medical spending on at-risk newborns through subsequent academic achievement](#).

PORTFOLIO OF RELATED WORK

Health Insurance: The Massachusetts Health Reform and The Affordable Care Act

My first paper on the Massachusetts reform with Jonathan Kolstad is important because it provided some of the first empirical evidence on the impact of the reform (Kolstad and Kowalski, 2012, [“The Impact of Health Care Reform on Hospital and Preventive Care: Evidence from Massachusetts”](#) *Journal of Public Economics*). In our paper, we examine welfare-relevant tradeoffs in a systematic way by comparing impacts on several measures of cost, quality, and access from hospital and survey data. Overall, we find that access and quality increased while costs grew in line with their previous trajectory. Our findings have inspired many subsequent studies in economics and medicine, and they have [received attention in several outlets](#). I have presented findings from my portfolio of work on the Massachusetts reform to over 50 audiences of academics and policymakers, including the Council of Economic Advisers, the Office of the Assistant Secretary for Planning and Evaluation, the Bureau of Economic Analysis, and the Census Bureau. Three of our most notable findings set the stage for future work in my portfolio.

First, we found that preventable admissions to the hospital decreased, but only when we controlled for patient severity. This finding is consistent with adverse selection into insurance before the reform such that healthier patients sought hospital care after the reform. In a short subsequent paper with Martin Hackmann, we found direct evidence of adverse selection into insurance using hospital and survey data (Hackmann, Kolstad, and Kowalski, 2012, [“Health Reform, Health Insurance, and Selection: Estimating Selection into Health Insurance Using the Massachusetts Reform,”](#) *AER Papers and Proceedings*). That study motivated our [later joint work](#) (discussed above), which quantified the welfare impact of adverse selection using data on the individual health insurance market.

Within a year of the establishment of the individual mandate penalty under the national reform in 2014, I extended our previous work on adverse selection in Massachusetts to analyze its impact on the individual health insurance markets in other states (Kowalski, 2014, [“The Early Impact of the Affordable Care Act, State by State,”](#) *Brookings Papers on Economic Activity*). Even though the ACA was a national policy, I found that state-level policies that affected its implementation had an economically significant welfare impact. In states that left enforcement to the federal government and states that did not set up exchanges, health insurance was less affordable because higher-cost individuals selected into the market. My findings appeared in several [outlets in the popular press](#).

Second, approximately half of new health insurance coverage was employer-sponsored despite fears that employers would drop coverage and instead pay the penalty established by the employer mandate. We explain this finding in our subsequent work on the labor market impact of health reform (Kolstad and Kowalski, 2016, [“Mandate-based Health Reform and the Labor Market: Evidence from the Massachusetts Reform”](#) *Journal of Health Economics*). The key to our advance is the link between the reform and our model, which incorporates impacts on the labor market through the individual mandate, the employer mandate, and subsidies for coverage outside of employment. Most Americans get health insurance through employers. [Previous theory](#) shows that employer-sponsored benefits distort the labor market by making it more expensive to hire

employees, decreasing labor demand. However, if employees value those benefits, they are willing to work for lower monetary wages, so labor supply also increases, decreasing or eliminating the distortion. Our extensions allow us to recover the size of the distortion.

Like our work on adverse selection, our work on the labor market has a tight link between the reforms, the model, and the estimation, and we depict the size of the distortion in a simple figure. Using data that follow individuals over time, we find that employees in Massachusetts after the reform valued health insurance at almost the full cost to employers. The distortion to the labor market is only 8% as large as it would have been if employers had been taxed to provide insurance that employees did not value. Our model and findings explain why employer-sponsored coverage increased despite fears that employers would drop coverage and pay the penalty: the individual mandate made employer-sponsored coverage more valuable. Employer-sponsored coverage was already valuable relative to other types of coverage because employers and employees pay premiums before taxes, and employer-sponsored coverage is often more generous and less expensive. For higher income workers not eligible for subsidies, the individual mandate tightened the link between health insurance and employment by making employer-sponsored coverage even more valuable. Our findings have received attention [in the popular press](#).

Third, hospital admissions from the emergency room decreased even though there was fear that emergency room usage would increase when coverage increased. The state of Oregon expanded Medicaid through a lottery in 2008, and findings from the resulting experiment got attention in the [New York Times](#) for showing that emergency room usage increased, in contrast to our finding and related work on the Massachusetts health reform [by other researchers](#). In recent work, I reconcile both findings using a model of selection into health insurance coverage that incorporates important features of the Oregon experiment and the Massachusetts reform (Kowalski, forthcoming, [“Reconciling Seemingly Contradictory Results from the Oregon Health Insurance Experiment and the Massachusetts Health Reform,”](#) *Review of Economics and Statistics*). Within the Oregon experiment, I find that sicker people who used the emergency room most when uninsured are more likely to sign up for health insurance coverage, and they increase their utilization the most upon gaining coverage. This finding can reconcile the results because Oregon expanded coverage to sicker people who signed up for a lottery, while Massachusetts expanded coverage to healthier people who avoided paying the penalty associated with the individual mandate.

The key advance that enabled my findings is that I saw fundamental parallels between models of adverse selection and econometric models used to examine [heterogeneous treatment effects](#). These parallels allow me to find adverse selection into health insurance coverage within the Oregon experiment and relate it to adverse selection within the Massachusetts reform. More importantly, they allow me to find that the treatment effect of health insurance coverage on emergency room utilization varies with selection into health insurance coverage. I originally demonstrated the parallels and explained the models using simple figures in a working paper (Kowalski, 2016, [“Doing More When You’re Running LATE: Applying Marginal Treatment Effect Methods to Examine Treatment Effect Heterogeneity in Experiments,”](#) *NBER Working Paper*). I subsequently divided the econometric content from the working paper between my work on the Oregon experiment and a separate paper on mammograms that I discuss with my recent and current work below. [Press coverage](#) of the working paper helped me to distill my findings. I have released accompanying [Stata commands](#) to make computation more accessible. In a third paper that does not break new ground, I use stylized examples to explain how the model and figures can be useful to examine external validity within experiments (Kowalski, forthcoming, [“How to Examine External Validity Within an Experiment,”](#) *Journal of Economics and Management Strategy*). To

better understand the implementation of experiments, I worked on a trial with a team of colleagues experienced in running trials (Anderson, Horn, Karlan, Kowalski, Sindelar, and Zinman 2021, [“Evaluation of Combined Financial Incentives and Deposit Contract Intervention for Smoking Cessation: A Randomized Controlled Trial,” *Journal of Smoking Cessation*](#)).

Health Insurance: The Tradeoff Between Price Sensitivity and Risk Protection

High deductible health insurance policies have become dramatically more popular since the Medicare Modernization Act of 2003 established health savings accounts, which allow people with high deductible health insurance policies to pay for qualified health expenses with pre-tax dollars. With two economists who worked with me at the Council of Economic Advisers in 2003 before I enrolled in graduate school, I examined data that we obtained from health insurers who sold high deductible plans on the individual health insurance market (Kowalski, Congdon, and Showalter, 2008, [“State Health Insurance Regulations and the Price of High Deductible Policies,” *Forum for Health Economics and Policy*](#)). Our finding that policies were prohibitively expensive or unavailable in states with the most restrictive regulations motivated my interest in the welfare impact of adverse selection in my later work.

The promise of high deductible policies to limit spending through consumer price-sensitivity inspired my dissertation research. In my job market paper, I found that people enrolled in employer-sponsored plans respond even more to the prices that they face for medical care than [classic results](#) from the Rand Health Insurance Experiment would suggest, in part because they face greater out of pocket costs than participants enticed to participate in an experiment (Kowalski 2016, [“Censored Quantile Instrumental Variable Estimates of the Price Elasticity of Expenditure on Medical Care,” *Journal of Business and Economic Statistics*](#)). To address the issue that some people had very high expenditures and many had no expenditures at all, I developed a new estimator with econometrician coauthors (Chernozhukov, Kowalski, and Fernandez-Val, 2015, [“Quantile Regression with Censoring and Endogeneity,” *Journal of Econometrics*](#)). We released a [Stata command](#) and accompanying [Stata Journal article](#) with Sukjin Han. My work was awarded the [Zellner Thesis Award in Econometrics and Statistics](#), and it also influenced projections by the Congressional Budget Office.

Price-sensitivity increases welfare by curtailing over-consumption of insured medical expenses, but it decreases welfare by exposing people with insurance to greater risk. In the remaining chapter of my dissertation, I developed a structural model to examine both sides of the tradeoff simultaneously by extending [classic theory](#) of responses to nonlinear taxation (Kowalski 2015, [“Estimating the Tradeoff Between Risk Protection and Moral Hazard with a Nonlinear Budget Set Model of Health Insurance,” *International Journal of Industrial Organization*](#)). Among enrollees in employer-sponsored health insurance plans, I found a net welfare loss from over-consumption of medical care because all plans provided ample protection against catastrophic risk.

Risk protection from high medical expenses is likely the most valuable to people who experience large health shocks and lose access to employer-sponsored health insurance. However, few sources of data follow people who change coverage or become uninsured. By securing restricted longitudinal data on all people who visited hospitals in the state of New York over a sixteen-year period, I found that Medicaid provides valuable risk protection to young people with private health insurance because they enroll in Medicaid after experiencing health shocks (Kowalski, 2015, [“What Do Longitudinal Data on Millions of Hospital Visits Tell Us About Public Health Insurance](#)

[as a Safety Net for the Young and Privately Insured?” *NBER Working Paper*](#)). I am significantly revising this paper with Kurt Lavetti and Lee Lockwood using richer data and theory.

Health Spending: The Role of Politics

Economists often consider [technology growth](#) to be the main driver of health spending growth. I am working with Zack Cooper and two political scientists to explore the role of a different driver: politics (Cooper, Kowalski, Powell, and Wu, 2020, [“Politics and Health Care Spending in the United States,” *NBER Working Paper*](#)). We find that politicians that vote to increase health spending receive larger hospital payments in their districts. In turn, local health spending increases and politicians secure larger campaign contributions. Our work has received attention in [the popular press](#).

RECENT AND CURRENT WORK

Health Care: Targeting Mammograms

Mammograms are controversial because they involve a tradeoff. The benefits of mammograms and the costs of “false positive” diagnoses are intuitive. Mammogram recommendations have weakened worldwide in response to growing evidence of costs that occur through a channel that is less intuitive: overdiagnosis of “true positive” breast cancers that would not eventually cause symptoms. No one can ever tell if a given diagnosis is an overdiagnosis, but long-term data from [clinical trials](#) show that overdiagnosis occurs. Women in the control arms only received mammograms during the trials if they experienced symptoms. Data from decades later show that smaller fractions of women in the control arms have ever been diagnosed with breast cancer, demonstrating that women in the intervention arms were overdiagnosed with breast cancers that would not have caused symptoms for decades.

Diagnosis of breast cancer can be costly because most women diagnosed with breast cancer pursue therapies such as surgery, chemotherapy, and radiation, which can cause side effects and even death. As I discuss in a symposium on [preventive care](#), clinical trials on mammography do not show statistically significant reductions in long-term all-cause mortality for women in any age group (Kowalski, 2021, [“Mammograms and Mortality: How Has the Evidence Evolved?” *Journal of Economic Perspectives*](#)). Given mounting evidence of overdiagnosis, the United States Preventive Task Force mammography guidelines for women in their 40s now recommend that women consult with their doctors and receive mammograms as they see fit.

Do these mammography guidelines target women who benefit most? I address this question using data shared with me by the investigators of the Canadian National Breast Screening Study, a large trial influential to the guidelines (Kowalski, 2023, [“Behavior Within a Clinical Trial and Implications for Mammography Guidelines,” *Review of Economic Studies*](#)). The data follow participants long enough to demonstrate overdiagnosis through examination of breast cancer as a health outcome. They also allow me to observe mammography behavior: some women in the control arm receive mammograms and some women in the intervention arm do not. Combining the data with a model of mammography behavior, I find that women more likely to receive mammograms are healthier and more likely to be overdiagnosed by them. Extrapolating my findings to the current environment using the model and findings from the literature, my findings suggest that further weakening of mammography guidelines could provide a rare opportunity to

improve health and decrease health spending. It would also make guidelines in the United States more consistent with guidelines in other countries.

The key to my advance is the connection between the clinical trial data and the model. Previous evidence from [Korea](#) and [the United States under current guidelines](#) corroborates my finding of selection into mammography such that healthier women are more likely to receive mammograms, but it does not consider treatment effect heterogeneity, and it does not consider overdiagnosis because it cannot observe long-term outcomes of women who do not receive mammograms. Using evidence from the trial data to motivate assumptions that I present graphically, I identify treatment effect heterogeneity using weaker assumptions than my previous work on the [Oregon experiment](#), discussed above. I also present my assumptions more simply. My work has the potential to be useful beyond the context of mammography because it illustrates how to use behavior within the same clinical trial data used to develop guidelines to inform targeting within those guidelines. My paper will be awarded the [Willard G. Manning Memorial Award for the Best Research in Health Econometrics](#) in June 2023.

Health Care: Targeting Medical Treatments that Save Some and Kill Others

Motivated by my work on mammograms, I find it interesting that surgery, chemotherapy, and many other medical treatments can save the lives of some people but kill others. The model that I use to examine treatment effect heterogeneity in my previous work allows random assignment to affect treatment by different magnitudes but in weakly the same direction. In my current work, I relax the common assumption that random assignment cannot affect people in opposite directions (Kowalski, 2020, [“Counting Defiers: Examples from Health Care,” ArXiv Working Paper](#)).

I make progress by building a model of the randomization process within an experiment, inspired by my work with previous models of treatment effect heterogeneity, which showed me that some parameters can be calculated in two ways that only differ because of the realized randomization process. An attractive feature of the model that I develop is that the key assumption depends on the randomization process: for example, a series of coin flips. The randomization process is part of the experimental design, so assumptions based on it can be made arbitrarily compelling through careful implementation. I demonstrate using simulations from the model that it is possible to infer with 95% confidence that at least three individuals are killed by an intervention that saves lives on average within a sample of 100 people. My results seem to contradict previous impossibility results. I show that those results rely on different data and assumptions.

This research is still in an early stage. I am working to understand how inference is possible even though the parameters are not identified in a traditional sense. The key advances seem to be the model of the data generating process, the focus on a finite sample as opposed to an infinite population, and advances in computational power that were not available to [early statisticians](#) who proposed simplifying asymptotic assumptions. I am also working to make computation more efficient and to incorporate more data to enable inference in a wide variety of applications.

Health Care: Targeting Covid-19 Vaccines: Implications for Equality vs. Equity

Targeting of medical treatments has implications for equality in access as well as equity in outcomes. In January 2021, Michigan Medicine offered appointments for Covid-19 vaccines through a series of randomized lotteries that continued through April, first to all patients aged 65 and older, and then within subsets of the original pool targeted by social vulnerability, race, and

ethnicity. I am currently working with David Chan and colleagues John Ayanian, Sarah Burgard, Sandro Cinti, James Henderson, Rahul Ladhania, Fiona Linn, Emily Martin, and Abram Wagner, from the University of Michigan Medical School, Michigan Medicine, and the University of Michigan School of Public Health to use the randomized design and a model to study how vaccination takeup varied by social vulnerability, race, ethnicity, and health. There have been attempts to measure Covid-19 vaccine equity, but they conflate access with vaccination takeup behavior. The multiple rounds of randomization will allow us to separate the two using weaker assumptions than previous models, and they will allow us to develop lessons to inform the targeting of health care in an important real-world context.