

Research Narrative

I am a development economist: **I study economic and social problems in the world's poorest countries.** In my research, I design and test practical solutions to real-world problems, while also shedding light on the economic factors that drive persistent poverty more generally. I do this by combining randomized field experiments with cutting-edge statistical and econometric techniques and economic theory. By running randomized experiments, I can answer causal questions like “how does this educational intervention affect learning?” More sophisticated experiments, and the use of economic theory, allow me to delve into the mechanisms behind my results. For example, in one study my coauthors and I tested whether the effectiveness of a savings intervention was driven by the fact that it serves as a commitment device, or because the savings are paid out in a lump sum. And we can learn even more from the data by using economic theory to understand the mechanisms behind the results, in order to draw conclusions that should hold in other settings. For example, in my work on HIV and risk-taking, I show that the data are consistent with a model that can also be applied to COVID-19.

This is all part of **an effort to make economics into an experimental science rather than an observational one.** The traditional view within economics is that this cannot be done, but I am part of an increasingly successful scientific movement to find ways to test economic theories using the randomized experiments—of the sort normally seen in the physical sciences.

In this document, I first summarize my overall productivity as a researcher. I then outline my main areas of active and published research, followed by my plans for future work.

Summary of Research Productivity

I am a highly productive development economist. **I have published 11 peer-reviewed articles**, including seven in economics journals. According to Google Scholar my work has been cited 377 times, and I have an i10-index of 12 and an h-index of 11. My published research includes an article in the flagship journal of the economics discipline, the *American Economic Review*. I have also published papers in the best journals in the fields of development economics (the *Journal of Development Economics*) and econometrics (the *Journal of Econometrics*), as well as the leading journal in demography (*Demography*). I played a major role in all aspects of these research projects: economics papers do not typically distinguish between lead authors and other roles, and my work adheres to that norm. With one exception noted on my CV, all of my economics publications are jointly lead-authored.

I have presented my work at **64 invited seminars and conference presentations**, including in leading forums such as the National Bureau of Economic Research Summer Institute (NBER SI). I also have six active working papers that are at the stage of completed drafts, and an active pipeline of new research projects at various earlier stages. To fund my research agenda, I have secured **ten external grants totaling nearly \$1.1 million**, as well as **\$159,000 in internal grant funding** from my own institutions.

Main Areas of Research

My research agenda thus far has focused on three of the most important problems facing people in developing countries today. One of my lines of research is focused on helping some of the world's poorest people save money—which they must do in order to make major purchases and insure themselves against negative shocks. In a second line of work, I study how to address the “learning crisis” in the developing world, wherein school enrollments have skyrocketed but students learn very little once they are in school. My third major line of research is about how people make decisions in the face of the worst pandemic disease of our lifetimes—HIV.

Helping the poor save money

Almost half of people in the developing world do not have access to a bank account. At first blush, this might not seem to matter much: given their low incomes, surely the poor do not have the capacity or need to save money? Yet there is strong evidence that not only do people have many reasons to save money, they make extensive use of a wide range of different savings methods. The poor need to save because their income is volatile—living on a dollar a day is more likely to mean an income of \$10 every ten days than a smooth, low daily income. They also need to save because their access to credit is limited, so they need to accumulate lump sums to make major purchases, such as paying school fees or buying major assets like improved metal roofs. Improved savings methods and higher savings rates could even be a pathway out of poverty: smoothing out one's income allows for better planning, and assets can either contribute to a business or save people money. For example, an improved roof is one that you do not have to pay to repair every year. Despite the importance of saving for the global poor, however, their savings rates are low, and they largely rely on high-risk methods like rotating savings groups and hiding money at home.

To address this problem, my coauthors and I study **a simple savings method: paying people later**. If a person earns a regular income from a job, their employer can create a low-tech savings method by giving them the option of receiving some of their pay in a later lump sum. My first paper on this topic, Brune and Kerwin 2019 (published in the *Journal of Development Economics*), showed that changing the time structure of workers' pay in this way led to changes in purchases of an artificial asset created for this study. My coauthor and I also documented a high rate of stated preference for getting paid later, which suggests workers valued this approach. A related paper (Brune, Kerwin, and Li 2022, published in the *World Bank Economic Review*), shows that exposure to a tempting environment on payday did not lead to increases in temptation spending. This paper includes my former PhD student Qingxiao Li as a coauthor, and exemplifies my apprenticeship-style approach to training doctoral students. I frequently involve my PhD students in research projects as a way to teach them the process of doing academic research.

Building on this work, my coauthors and I **partnered with a large agricultural firm to offer deferred wages to its employees**. Workers were allowed to choose whether to participate in the scheme, and how much to have withheld from each paycheck; all the savings were paid out in a single lump sum at the end of the harvest season. My coauthors and I won a \$344,975 grant from the Innovations for Poverty Action (IPA) Financial Inclusion Program Research Fund to implement this study.

This savings scheme was extremely popular, and led to increases in asset ownership: two years later, the rate of improved roofing was ten percent higher in the treatment group than in the control group. This is one of the few examples of a savings product that actually leads to long-run increases in wealth—a major advance toward finding ways to help the poor meet their savings goals. This paper, Brune, Chyn, and Kerwin (2021), was recently published in the *American Economic Review*. In ongoing work, we are exploring ways to get employers, NGOs, and governments to give payment recipients the option of getting paid later as a way to help them save.

With the same coauthor team, I also used data on worker output from the tea company to **study peer effects at the workplace**—showing that they are likely driven by motivation, rather than shame, and thus could be used to improve worker welfare. This paper is forthcoming at the *Journal of Human Resources*. We also have ongoing work using this data to estimate workers’ time preferences and labor supply elasticities via bunching methods from the public finance literature.

Addressing the learning crisis

“No child left behind” has been a mantra in US education for nearly two decades now, and while the specific program has many critics, nobody disagrees with the principle behind the name. But in developing-country education systems, many children are left behind, or never go anywhere in the first place. In light of this, my coauthors and I have been collaborating with an educational organization in Uganda for the past nine years to study the Northern Uganda Literacy Project (NULP). The NULP is a mother-tongue-first literacy intervention aimed at the first three grades of school in the Lango sub-region, a formerly conflict-affected part of northern Uganda. Our research focuses on measuring how well the program works to improve learning, why it works, and its long-run consequences for students.

The study of the NULP is a panel randomized trial in which schools were assigned to one of three study arms: a control group, the original NULP, and a reduced-cost version of the NULP that was designed to simulate how it might be scaled up. Our initial findings showed that the NULP raised test scores substantially, putting students nearly a year ahead in terms of reading ability by the end of first grade, with substantial benefits for writing scores as well. This makes it **one of the most effective education interventions in history**. However, the reduced-cost version of the program had more mixed results: we found that it raised basic reading skills and writing skills, but actually backfired for advanced writing ability, leaving students with *lower* writing ability than the control group.

To understand this finding, we **developed an economic model that yielded two possible explanations**. First, the different inputs the model provides could be strongly complementary to one another: a lack of one could reduce the impact of another. Second, the effects of the program could follow a “J-curve”—student learning could dip below zero before going up. We tested these using machine-learning methods and by drawing on detailed information about the intervention and teaching practices in the schools, finding evidence for both explanations. This initial work, Kerwin and Thornton 2021, was recently published in the *Review of Economics and Statistics*.

In a follow-up paper that is forthcoming at the *Journal of Econometrics* (Buhl-Wiggers et al. 2022), we examine how the treatment effects of the program vary across students, demonstrating that **even this highly effective program continues to leave many students behind**. In another

working paper with many of the same coauthors, we use the data from this study to examine how much teachers in Uganda vary in their effectiveness at improving learning, and how the NULP program changes that.

Our ongoing work is using the NULP data to study additional questions: How did actually scaling up the program play out? What are the long-run effects of improved literacy on learning in other areas, and on eventual labor market outcomes? We have received a \$148,647 grant from the Abdul Latif Jameel Poverty Action Lab (J-PAL) Post-Primary Education fund to study this last question, by continuing to track the original cohort of students as they proceed into secondary school. We are currently collecting data for an **eight-year follow-up of the original intervention**.

In a related line of work, I am working with a separate team of coauthors to study the effects of a remedial learning program in Odisha, India. This intervention also is meant to address the learning crisis, but at a different stage: many students in India reach secondary school while still lacking elementary-school-level skills. The program, called *Utkarsh*, provides students with targeted remedial lessons to bring them back up to their grade level. We have secured three external grants worth a total of over \$400,000 to fund this research. Our preliminary findings show that the program raised test scores throughout the distribution of (initial) student learning levels, benefiting not just the weakest students but also the strongest ones. As part of the study, we also tested an alternative approach to the program that provided teachers with more flexibility in how much remedial instruction they did. We find very little difference in the program's effectiveness, and in fact show that **teachers have little demand for this additional flexibility**. We recently completed data collection for a long-run follow-up that will examine how the program affected advancement to the next level of secondary school.

Decisionmaking and pandemic disease risks

Economists and public health scholars have long known that health risks lead to “risk compensation”: when an activity becomes more dangerous, people tend to do less of it. For example, the United States saw widespread declines in restaurant reservations in March of 2020, even in places with no formal lockdowns or bans on indoor dining—because the increased risk made people less willing to eat out. Given this phenomenon, one potential strategy is to exaggerate risks. If risk compensation already leads people to avoid risky activities, telling them that activities are extremely risky will be even better, right? Wrong. In my work on HIV and risk compensation, **I show that this “scared straight” style of messaging can backfire**, leading people to become fatalistic and take more risks rather than fewer. The reason this happens is because a higher risk from unprotected sex makes additional sexual activity more dangerous—but also makes it more likely that you already have HIV, in which case it no longer matters what you do. I test this theory using a randomized experiment in which I taught people about the true risk of HIV transmission from unprotected sex, which is substantially less than most people think. Conventional risk compensation predicts that this new information will lead people to have more sex. Consistent with the model, however, people with high initial risk beliefs react to this information by having *less* sex—exactly as we would expect to happen if people's exaggerated risk beliefs made them fatalistic. This paper, Kerwin 2022, is currently under review. It won the 2016 Dorothy S. Thomas Award from the Population Association of America, was featured on NPR's *Morning Edition* in 2016.

A related paper using the same dataset shows that **when interviewers know more about the risk belief questions they are asking on a survey, their knowledge spills over onto the recorded risk beliefs collected by the survey.** This is an important consideration in the design of surveys to capture subjective risk beliefs, which are increasingly common in social science. This study, Kerwin and Ordaz Reynoso 2021, was published in *Demography* and is coauthored with my former PhD student Natalia Ordaz Reynoso. Together with Divya Pandey, another one of my PhD students, I am also using the data from this experiment to explore “epistemic uncertainty”, which captures how certain (or uncertain) people are about their risk beliefs. Our results show that epistemic uncertainty can be measured separately from levels of risk beliefs, and is predictive of how much people update their beliefs in response to new information. We have already presented these results at an external seminar, and are planning to complete a draft of the paper this summer.

Building on this previous work on HIV, I am now collaborating with a team of other researchers to explore how to ensure that people who are already infected with the virus get tested, which will give them access to life-saving treatment. Based on insights from behavioral economics, we designed multiple interventions to increase the take-up of HIV testing in southern Malawi, where nearly one in five adults has the virus. One intervention was a traditional financial commitment device: people could choose to stake part of their payment for participating in the study on getting an HIV test. This choice then committed them to get tested, since otherwise they would not receive the money. The other intervention was an appointment for a test, which we show can function as a highly effective substitute for a commitment device. Using a randomized experiment, my coauthors and I found that the appointments worked much better than the commitment devices, while both outperformed a control group. Moreover, we show that the effects of appointments are concentrated among men who wanted to enroll in the commitment device, while men who signed up for commitment devices commonly failed to follow through and thus lost the money they staked on getting an HIV test. **Appointments address self-control problems even better than conventional commitment devices,** and without the downside of potentially costing people money.

Since the COVID-19 pandemic began in 2020, I have also extended my work on decisionmaking about disease risks to this novel threat to population health. In a study published in the *Journal of Economic Behavior and Organization* (Fitzpatrick et al. 2021), my coauthors and I show that knowledge of the symptoms and transmission mechanisms of COVID-19 was high across a sample of people in four countries in Africa. However, **we find fairly low correlations between knowledge and protective measures—and show that greater knowledge was associated with less social distancing rather than more.** The severity of the pandemic in the United States has also motivated me to apply my expertise to mitigating the disease’s effects here. Together with Marc Bellemare and Kent Horsager, I am conducting a pilot study of a simple method to improve ventilation in schools in Minnesota. Our approach uses CO₂ monitors to detect dangerously poor ventilation, since CO₂ is highly correlated with COVID-19 transmission risks. When the sensor indicates high CO₂ levels, teachers are instructed to open a window for 15 minutes. Our initial data suggests this system is effective at improving air quality in schools.

Future Research Directions

In the longer term, I hope to focus my research on three key areas that build on my existing research portfolio and also will help to shape the broader field of development economics.

Scaling effective interventions

How do we get policymakers to actually implement interventions that work? International development is facing what might be termed a “scaleup crisis”, in which many programs are effective at a small scale, but the results do not replicate at large scale. One key driver, which Rebecca Thornton and I explored in our *Review of Economics and Statistics* paper, is that the scaled version of the program often changes elements of the original—and it is not always obvious how important the changes will be. Thus one a potential solution to the scaleup crisis is to **convince governments, NGOs, and so forth to implement programs as they are originally designed**, rather than changing them. Some recent research has begun to address this issue,¹ but there is still much that we do not know about how to achieve this.

I have begun to work on this issue from two perspectives. First, I am working on finding ways to scale up effective interventions that I have studied. Along with my coauthors Lasse Brune and Eric Chyn, I am exploring ways to implement the deferred income payment scheme from our *American Economic Review* paper as a savings tool. We are collaborating with GiveDirectly, which provides direct cash transfers to the poor, to provide this as an option for their transfer recipients. Second, I am looking for opportunities to work on other projects that are scaling up effective interventions. For example, I was recently contacted by an organization that is looking to scale up “graduation” programs, which provide poor people with a bundle of services including productive assets, cash grants, access to savings, and training. I hope to pursue this project or one like it in the future as a platform for studying and understanding how to maintain fidelity to programs at scale.

Understanding treatment effect heterogeneity

Empirical microeconomics focuses largely on the average effects of programs and policies. But treatment effects can vary widely from person to person. Another line of future work I plan to pursue is **estimating this variation for more development interventions** and developing insights into why treatment effects vary. This is a major shortcoming of the international development literature. For example, despite hundreds of randomized experiments studying the effects of education interventions, our “Some Children Left Behind” paper (Buhl-Wiggers et al. 2022, *Journal of Econometrics*) is the first to estimate the non-parametric Fréchet-Höfdding bounds on the variance of the treatment effects. One obvious next step that my coauthors from that study and I hope to take is to fill that hole, by conducting a meta-analysis of the Fréchet-Höfdding bounds for the education literature in developing countries.

Another direction that I hope to take in my future work is to **learn more about why treatment effects differ from person to person**. A key reason this is important is that it is another explanation for why programs do not work at scale, beyond changes in the program. If we could replicate the exact treatment again, and could forecast everyone’s individual treatment effects

¹ See, for example, Hjort et al. 2021 (*American Economic Review*) and Mehmood et al. 2021 (Working Paper).

(which could potentially vary over time and depend on who else receives the treatment), we would know exactly how well a program would scale up. To do this, we would need to not just describe how treatment effects vary but also be able to attribute that variation to specific known factors. There is some exciting new research that gets close to this goal in certain domains.² On the other hand, for education in particular, things are less promising: even modern machine-learning methods could explain very little of the variation in treatment effects that we document in “Some Children Left Behind.” One avenue for doing better here is to use economic theory to generate better insights into why treatment effects vary. Doing this has been an ongoing theme in my research: my “Scared Straight” paper about fatalism and HIV risks is fundamentally about exactly this approach. I hope to apply this to other settings, particularly education, going forward.

Belief miscalibration

Economic decisions depend fundamentally on people’s beliefs about the current and future state of the world. What will labor market conditions be like next year? How dangerous is eating out at a restaurant today? A common thread in these beliefs—exemplified in my research on responses to disease risks—is that people’s perceptions deviate sharply from reality. For example, the median person thinks that HIV is hundreds of times more transmissible than it actually is. Where do these miscalibrated beliefs come from? How do people actually form and update their beliefs? Is this related to basic misunderstandings about how probabilities work? I hope to continue my research agenda on understanding responses to disease risks by **focusing on the sources of these gaps in people’s knowledge and beliefs**. My ongoing work on epistemic uncertainty provides one way of understanding this issue: people seem to hold binary risk beliefs with varying degrees of certainty. I also hope to pursue other projects in this area that examine the role of misunderstandings about probability and the information sources that people draw on to form their risk beliefs.

² For example, Meager (2022, *American Economic Review*) shows that the positive effect of microcredit on the top of the distribution of profits is driven by households with previous business experience—which is consistent with previous theoretical and empirical work on the topic.