Statement of Research, Teaching, and Service • Daniel Björkegren • February 2026

Research

I am a development economist who analyzes the causes and consequences of the spread of digital technologies in low-income countries. Digital technologies raise new policy questions, enable new measurement, and can deliver new services. I work on each of these facets with a variety of tools, including structural modeling techniques from industrial organization, randomized experiments, and machine learning.

While my work is inspired by the lives of the world's poor, I try to find principles that are relevant for all. I partner with firms, nonprofits, and governments to obtain rich, digital data on populations that are sparsely documented in survey and traditional data.

> A first strand of my work analyzes the industrial organization of vital networks in low-income countries. Mobile phones have spread dramatically, knitting even the poor into the global economy. My work analyzes the value that mobile phone networks provide, what drives consumers to adopt, and how the industry can be regulated. I also consider an important transportation network in emerging cities: private transit.
>
> The spread of computation has enabled decisionmaking to be automated, which I explore in a second strand work on applied machine learning and artificial intelligence. I explore how such algorithms can be made more robust and transparent, and the potential of AI to improve service delivery in low resource settings. My work has provided foundations for one such service: digital credit scoring for people excluded from traditional banking. Follow up work has assessed the welfare impacts of digital loans. Additional work develops methods to align policies with the values of constituents.
>
> A third strand of my work explores how digital data can measure human behavior in detail new to the social sciences.

1. Industrial Organization of Networks

1a. New Mobile Phone Networks and Direct Network Effects

Mobile phone networks increasingly act as gatekeepers to information and finance across developing economies. Yet, there is little work on adoption, the supply of service, or regulation. How should societies tax networks, ensure coverage to remote areas, and ensure sufficient adoption? Should societies tolerate monopoly networks, or split them up?

A central challenge in answering these policy questions is that phone networks produce direct network effects. As more of the people you wish to interact with adopt phones, it becomes more beneficial for you to adopt. Because of these linkages, a policy that affects one person's adoption can cause ripple effects through the entire network. These effects are similar to those that characterize other important tech platforms (such as social networks, messaging platforms, and payment systems). Because developing country networks are still growing, these effects are more important than in mature phone networks.

These ripple effects make it difficult to analyze policy in any direct network industry. First, it is difficult to gather data on the entire network over which effects may spread. Second, even with this data, it is difficult to learn how one person's decision will affect another's. Two individuals who are connected may adopt at similar times because they value interacting with each other, or because connected people are exposed to similar shocks. Because phones are durable it matters not only which contacts have a phone at a given point in time but also who is likely to purchase a phone in the future, which may depend in turn on who whether you purchase a phone in the first place. There thus tend to be multiple equilibria. There is an empirical literature analyzing policy in simpler *in*direct network industries, in which consumers benefit from additional users not because they value links with them, but because popular platforms are better served by the other side of the market (for example, a rider cares about how many Uber drivers there are, but not their identities). However, there is very little work using data to evaluate policy in industries with direct network effects.

The article, **The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda** (*Review of Economic Studies*, 2019), overcomes these challenges to estimate and simulate the adoption of nearly every mobile phone subscriber in Rwanda from 2005-2009, using 5.3 billion transaction records. Because phone calls were billed by the second, a subscriber must value each connection at least as much as the cost of calls placed across it. Further, because the firm changed calling prices and increased the quality of service, I can identify the underlying demand curve for communication across each link—and thus the value of each link in the network. I develop a simulation method that allows each consumer to adjust their adoption and usage in response to a policy and to each other, capturing how the effects of a policy ripple through the network and across physical space. In the model, if one person adopts earlier, that can only increase the incentives for others to adopt earlier; this supermodularity property allows me to characterize the entire set of multiple equilibria in adoption. This first article uses this method to produce three empirical results:

First, it provides the first micro-identified estimates of the net welfare generated by a developing country mobile phone network: the equivalent of 2-3% of Rwanda's GDP over this period, with 51% accruing to consumers, 35% to operators, and 14% to government.

Second, it finds that taxing the adoption and use of growing digital services can be extremely distortionary. The mobile industry contributed an average of 7% of government revenues in sub-Saharan Africa (GSMA, 2008). Standard measures used by industry would suggest that these taxes induce a welfare cost to Rwanda of $1.22 per dollar of government revenue, but accounting for network effects, welfare costs are 2.4 to 4.5 times larger.

Third, it finds large spillovers from serving rural areas. When the operator was required to build towers in remote regions, 78% of the consumer benefits accrued to individuals in *other* regions. This suggests that remote communities are unlikely to adequately fund their own coverage. Requiring the operator to provide service to remote areas boosted welfare.

Overall, I find that network effects are first order in determining the effects of policy.

The model in this article also provides the foundation for two additional articles.

Network industries may tend towards monopoly, which has led to public concern about dominant tech firms. How should societies discipline these industries? A single network may internalize more network benefits and thus invest more, but also may take advantage of consumers. But it is challenging to empirically analyze competing networks. In markets that are already competitive, it is difficult to obtain and link network data from all competitors. Each firm must choose a path of actions over time anticipating what each other will choose, anticipating what consumers will demand, where each consumer values the usage of potentially millions of other consumers.

The article, **Competition in Network Industries: Evidence from the Rwandan Mobile Phone Network** (*RAND Journal of Economics*, 2022) is to my knowledge the first analysis of competition between direct network goods that uses data on individuals. Building on the previous model, it evaluates what would have happened if a competitor network were allowed to enter. I combine data on nearly the entire network from the original firm, and data on an eventual competitor that the government allowed into the market later, from a consumer survey I fielded. The model allows two competing firms to decide prices and investments in coverage, and consumers to select and switch operators. It characterizes the space of multiple equilibria by exploiting supermodularity in two components. It finds that adding an additional competitor earlier could have reduced prices 30-50%, and *increased* incentives to invest in rural towers. Such a policy would have increased total social welfare by the equivalent of 1% of GDP. However, incentives to invest would be lowered if competition were pushed too strongly, by forcing firms to interoperate their networks without charge (a policy proposed in the U.S.).

Another problem in networks is attaining sufficient adoption. The article, **Network Adoption Subsidies: A Digital Evaluation of a Rural Mobile Phone Program in Rwanda**, (with Burak Ceyhun Karaca, *Journal of Development Economics*, 2022) analyses the effects of a government adoption subsidy. A common concern with targeted subsidies is leakage: intended recipients may transfer the good. However, most subsidized handsets remain in rural areas where the subsidy was targeted, and are used as much as handsets purchased at retail prices. Simulations suggest the enacted handout had a positive social impact, but alternate voucher policy could have had a higher impact.

Altogether this work suggests that mobile networks provide large benefits but that obtaining good societal outcomes requires considering network effects carefully. I develop some of the first tools to do so empirically in this industry. I find that revenue estimates can be biased ranging from 52% too small to 86% too large in demand systems that do not model the full structure of the network.

1b. Private Transit Networks

Transit is another important network industry. In many emerging cities, transit is private: operated by many small bus companies or individuals. In recent years, many cities have invested in public transit, but these systems are costly and unlikely to displace private transit anytime soon. A series of projects seek to understand private transit in Lagos, which has a population of 22 million and is projected to become the world's largest city. 75% of motorized trips in Lagos are via private minibuses.

The project, **Public and Private Transit: Evidence from Lagos** (with Alice Duhaut, Geetika Nagpal, and Nick Tsivanidis, revision requested, *American Economic Review*) finds that private minibuses reduce service when the city introduces 820 new public buses across 40 routes. This increases wait times for commuters, dampening the welfare benefits of the new system. We assess impacts using data we collected. We stationed enumerators to monitor 278 minibus routes at terminals and 79 bus stops, over 13 rounds, as the public system was rolled out. When the government enters a route, minibuses depart less frequently, driver profits fall, and drivers switch to connected routes, reducing prices. We assess how commuters value the change by developing a custom app that works even with basic phones, and running an experiment providing random offers to wait. Waiting is authenticated through cryptographic codes that change each minute. We point out an issue affecting many revealed preference estimates of the value of time: individuals may be differentially likely to engage with a waiting decision when they are busy. For example, recent value of time estimates from rideshare platforms will be inflated if travelers tend to travel via rideshare when short on time. We use randomized daily offers to correct for this selection in our experiment; we show that without this correction estimates would be biased by 56%. Based on these preferences, we find that the private response harms commuters on treated routes, who wait longer, but benefits those on connected routes, who face only lower prices. The disciplining effect of the new system on prices dominates on average, so that commuters overall benefit from the introduction of public transit, while minibus drivers lose revenue. Over one quarter of the commuter welfare gains of building the public transit system arise from the response of private transit.

2. Applied Machine Learning and Artificial Intelligence

I work on applications of machine learning and AI and the questions they raise.

2a. Digital Credit

Mobile phones can offer digital services at close to zero marginal cost. This makes it viable to serve some of the 1.7 billion people who lack access to traditional financial services.

Because many of these people lack formal financial histories, they lack credit scores. My paper, **Behavior Revealed in Mobile Phone Usage Predicts Credit Repayment** (with Darrell Grissen; proposal posted online 2010; working paper 2015; *World Bank Economic Review*, 2020) shows that nuances captured in the use of mobile phones themselves can be used to predict repayment for unbanked borrowers. It combines phone transactions, data on repayment of phone credit, and any records present in the credit bureau from a Latin American country. It finds that models trained on digital traces are not perfect, but achieve performance comparable to how well credit bureau models perform on thin files.

A companion paper, **The Potential of Digital Credit to Bank the Poor** (with Darrell Grissen; *American Economic Association Papers and Proceedings*, 2018) shows that digital platforms can theoretically expand access to credit, by lowering the fixed cost of serving consumers. I find that performance of digital credit scoring declines only modestly as digital footprints are made sparser, so thus may be applicable to lower income consumers.

Digital credit products using similar forms of credit scoring have seen dramatic uptake in developing countries. This has led to widespread concern that borrowers are being taken advantage of (Donovan and Park, 2019). The paper, **Welfare Impacts of Digital Credit: A Randomized Evaluation in Nigeria** (with Joshua Blumenstock, Omowunmi Folajimi-Senjobi, Jacqueline Mauro, and Suraj Nair; *Economic Development and Cultural Change*, 2025) provides the first experimental estimates of the welfare impact of digital credit in a developing country. It joins two quasiexperimental evaluations of digital credit in the literature (Bharadwaj et al., 2019; Brailovskaya et al., 2021). We ran a field experiment that randomized the approval and amount of thousands of loan offers in Nigeria, and later surveyed respondents to measure effects on welfare. Randomly being approved for credit increases subjective wellbeing, and we rule out large positive or negative effects on financial health, resilience, and women's economic empowerment. We find little evidence of behavioral traps: if anything, borrowers tend to overestimate future borrowing.

2b. Robustness and Manipulation

Manipulation-Robust Prediction (with Joshua Blumenstock and Samsun Knight; conditionally accepted (data), *American Economic Review*) acknowledges that machine learning models designed to describe the world can be poorly suited to making decisions. Once an algorithm is implemented, it can be gamed, which may result in decisions that are poor or unsafe. This problem arises because the standard estimators used to construct decision rules assume that the relationship between the outcome of interest and human behaviors is stable. But this assumption tends to be violated as soon as a decision rule is implemented, which generates incentives for agents to change their behavior (Lucas, 1976). In absence of a framework, tech firms and implementers apply ad hoc solutions, like repeatedly applying estimators that assume that behavior will be stable—which are perpetually mistaken. Building on a recent theoretical literature in computer science (Bruckner and Scheffer, 2011; Hardt et al., 2016), we embed a model of behavior within a common estimator, to create an estimable method that anticipates gaming. Our approach generates rules designed to remain robust when implemented and made transparent.

Typical datasets used in supervised machine learning observe behavior in a single environment, prior to the implementation of a decision rule, or with a single static rule. Manipulation is not detectable in these datasets, because implementing a decision rule changes the environment. We introduce the term 'counterfactual fit', to denote the loss in a counterfactual environment, such as where a decision rule is implemented and made transparent.

To estimate counterfactual fit, we develop a new platform for implementing different decision rules experimentally. We created a new app to mimic digital credit and ran a field experiment in Nairobi, Kenya with 1,557 people. The app gathers the same data as a digital credit app, and experimentally varies the decision rule that each individual faces each week, and whether they see the decision rule (transparent treatment) or just a description of what is being predicted (opaque). We use this setup to experimentally estimate the costs of manipulating different behaviors, for different types of individuals. Behaviors that appear similarly predictive to a standard machine learning method can be wildly differentially manipulable. We find that a 'strategy robust' approach that anticipates manipulation is more accurate than standard methods when implemented.

The paper also contributes to the policy debate on algorithmic transparency: a concern raised by tech firms is that if decision rules (like Google search, or Instagram's ranking) were disclosed, they would become so manipulated that they would cease to function. Our method can simulate the equilibrium change in performance if the details of a decision rule were disclosed. We find these costs are small in our context.

It also connects to a debate on the role of theory in machine learning. The field of machine learning has made substantial progress in recent years by combining atheoretical approaches with massive amounts of data. However, in applications where counterfactual fit is the appropriate measure of performance, one will rarely have sufficient data from the exact environment of interest. We show that a form of *structural machine learning* that models the theoretical relationship between the model and resulting incentives to manipulate can improve performance.

2e. Alignment of Policies and Values

A related strand of work aims to assess and align policies with preferences. This is relevant for many policies, and increasingly important as decision rules are encoded in algorithms.

While there is substantial work across fields on how to make decisions when an objective function is provided, it is difficult for societies to weigh different harms and benefits to provide such a function. **What do Policies Value?** (with Joshua Blumenstock and Samsun Knight; forthcoming, *Review of Economic Studies*; computer science article published nonarchival in ACM EAAMO, 2021) develops a method to infer the preferences over households and outcomes that are consistent with an allocation policy. It does so by comparing who benefits most from an intervention (using estimates of heterogeneous treatment effects), with who it is allocated to (based on eligibility criteria). We apply it to Mexico's PROGRESA anti-poverty program. While certain subgroups—such as indigenous households—were ranked highly in the allocation, those groups benefited so much more from receiving the program that the implicit welfare weight on them by the policy was lower. We find that the implied value of health and schooling are consistent with paternalism, and that the implemented policy was broadly consistent with preferences elicited from Mexican constituents through an online survey. Our method allows societies to invert policy discussions, to debate ends (welfare and objective weights) rather than their means (the concrete implementation details).

A related concern is that algorithms that are tuned for one objective can cause problematic side effects. **Balancing Competing Objectives with Noisy Data: Score-Based Classifiers for Welfare-Aware Machine Learning** (with Esther Rolf, Max Simchowitz, Sarah Dean, Lydia Liu, Moritz Hardt, and Joshua Blumenstock; *International Conference on Machine Learning (ICML)*, 2020) develops a theoretical framework for optimizing multiple objectives. We show that if a primary objective (e.g., welfare) can be measured only noisily, it may be better served by optimizing a secondary objective that is better measured if it is sufficiently correlated (e.g., engagement). We apply the framework to YouTube's recommendation system. The conference article is published in a top machine learning venue; the workshop version of the paper earned a best paper award at the *Neural Information Processing Systems (NeurIPS) Joint Workshop on AI for Social Good*, 2019.

The short paper, **Credit Scores that Prioritize Customer Welfare: Theory and Evidence from Nigeria** (with Simón Ramirez Amaya, Joshua Blumenstock, and Suraj Nair; *International Conference on Learning Representations (ICLR) Workshop on Practical Machine Learning for Developing Countries*, 2023), proposes a way for lenders to trade off repayment likelihood against benefit to borrowers using multi-objective machine learning.

2b. Methods for Causal Inference

Machine learning methods also provide new opportunities to address traditional econometric challenges. **Causal Inference from Hypothetical Evaluations** (with B. Douglas Bernheim, Jeffrey Naecker, and Michael Pollmann; rejected with an invitation to resubmit, *Journal of Political Economy*) develops a new method for inferring the causal effect of treatments on choices. A central econometric challenge is that real choices are only observed in actual treatment states. In place of real choices, much work analyzes hypothetical choices, which can be asked of any treatment state—but responses to hypothetical questions are known to be systematically biased. However, that these biases are systematic suggests that hypotheticals encode relevant information, and can make good predic*tors* of actual choices, even if they are bad predic*tions*. We combine these two sources of data to estimate the relationship between real choices and hypothetical responses. Hypothetical responses may include hypothetical choices as well as ratings of motivations that can influence the hypothetical bias (such as whether a choice is tempting or would be viewed positively by others). We use this estimated relationship to undo biases in hypothetical choices. Relative to other methods, ours accounts for factors causing differences in hypothetical bias across treatment states, by exploiting variation across a collection of similar choice settings. Under appropriate conditions, it can recover causal effects of treatments that have yet to be implemented, or treatments assigned endogenously, under conditions where quasiexperimental methods fail. Our approach yields estimates close to ground truth in microfinance and lab applications.

2d. Artificial Intelligence

Only 37% of Africans use the internet, and those who do rarely use the web. The most frequently cited reason for low internet usage is cost. The article **Could AI Leapfrog the Web? Evidence from Teachers in Sierra Leone** (with Jun Ho Choi, Divya Budihal, Dominic Sobhani, Oliver Garrod, and Paul Atherton, submitted) studies an AI chatbot for teachers in Sierra Leone that is accessible in a common instant messaging app. These teachers use AI more than they use web search. We compare the responses of AI to the web search results that would have been retrieved for the same queries from most popular local search engine. The average web page consumes 3,107 times more data than an AI response. Including the cost of compute, AI is 98% less expensive than loading a web page. Additionally, only 2% of results for corresponding web searches contain content from in country. We develop new text analysis techniques to understand the queries submitted. In blinded evaluations, an independent sample of teachers rate AI responses as more relevant, helpful, and correct than web search results. These findings suggest that AI-driven solutions may have the potential to cost-effectively bridge information gaps in low-connectivity regions.[1]

[1] An earlier workshop version of this article was presented at *NeurIPS Workshop on Generative AI for Education*, 2023 (**Are LLMs Useful in the Poorest Schools? TheTeacher.AI in Sierra Leone** with Jun Ho Choi, Oliver Garrod, Paul Atherton, Andrew Joyce-Gibbons, and Miriam Mason-Sesay).

There is widespread debate about how AI will affect the economy. A recent literature has suggested that interest rates provide a signal of what market participants believe may happen (Chow, Halperin, and Mazlish 2025). Andrews and Farboodi (2025) find that releases of proprietary AI models coincide with large decreases in yields of long-term bonds. My paper, **Market Beliefs about Open vs. Closed AI** (working paper), finds the effects are opposite around the release of open-weight AI models that can be freely run and modified, and that cumulative effects are close to zero. This suggests that the shifts in bond yields around AI models could result from other factors such as investment.

3. Big Data for Development

A common strand through my work is the use of digital data in developing economies.

Researchers and policymakers increasingly rely on measures of mobility from mobile phones. However, these measures may not be representative of population movement, particularly in low-income communities. The most widely used mobility data comes from smartphones. **Assessing Bias in Smartphone Mobility Estimates in Low Income Countries** (with Sveta Milusheva and Leonardo Viotti; *ACM Conference on Computing and Sustainable Societies (COMPASS)*, 2021) uses telecom data from Uganda, where few adults had smartphones. Smartphones reduce mobility much more than basic phones in response to a pandemic lockdown, so yield biased estimates of population mobility.

Another source of data comes from mobile phone operators, whose systems track the approximate locations of both smartphones and basic phones. However, the most common source (call detail records) records locations only during transactions, so that movement is endogenously missing. The paper, **Representation in Mobility Data in Emerging Cities** (with Alice Duhaut, Oluchi Mbonu, Geetika Nagpal, and Nick Tsivanidis; *American Economic Association Papers and Proceedings*, 2026) uses richer, continuous location data from a telecom covering subscribers at 15 minute intervals in Lagos, which has higher rates of smartphone adoption. It finds that both smartphone and transaction data capture residential population distributions with similar fidelity as the continuous data, but that both sources can distort measures of trips and transportation flows across the city.

Measuring Informal Work with Digital Traces: Mobile Payphone Operators in Rwanda (*International Conference on Information and Communication Technologies and Development (ICTD)*, 2020) creates a dataset of the minute-by-minute business decisions of informal vendors, and uses this to trace their learning curves as they enter the industry.

4. Ongoing work

I also have several ongoing projects:

Efficiency of Informal Transit Networks (with Alice Duhaut, Geetika Nagpal, Nicola Rosaia, and Nick Tsivanidis, in progress) assesses the degree to which competing minibus firms provide efficient service in Lagos with a structural model. We will assess if competing providers overenter on central routes, or provide too little service on the fringes.

Trained medical professionals are scarce in many low-income settings. In Kenya, just two doctors serve every 10,000 people, compared to 37 in the United States (WHO). As a result, patients often turn to pharmacies for front-line care. **Can AI Improve Frontline Care? Evidence from Retail Pharmacies in Kenya** (with Michael Dinerstein, Samantha Horn, and Anne Karing, in progress) will investigate how generative AI assistance interacts with consumer preferences and provider incentives.

Overall evidence of impact of research:

As overall measures of the impact of my research, I am regularly invited to give seminars at institutions such as Harvard, MIT, Stanford, Princeton, NYU, Yale, Dartmouth, UC Berkeley, the University of Chicago, and Washington University in St. Louis. These seminars have been in a variety of groups at economic departments (development, industrial organization, econometrics, finance, applied microeconomics), business schools, and computer science departments. I have presented at top conferences, including the *National Bureau of Economic Research* (NBER) 12 times and the *Bureau for Research and Economic Analysis of Development* (BREAD) 1 time.

With peers I have raised 12 federal and private foundation grants worth a combined total of over $1.5 million. Including grants with senior colleagues Michael Kremer and Shawn Cole, I have raised over $4.9 million. I am regularly invited to give keynote speeches at research conferences (7 times), as noted on my CV.

Broader Policy Impacts

My 2010 proposal for digital credit scoring was one of the first public documents proposing credit scoring based on digital traces. The paper itself was given a 4-minute segment on *NPR Morning Edition* in 2015. It was cited in *The Impact of Machine Learning on Economics* (Athey, 2018), as well as *The Age of Surveillance Capitalism* (Zuboff, 2019), which was rated as a top nonfiction book by *The New Yorker*. This approach is now commonly used for applicants who lack financial histories in a new multibillion dollar industry providing digital loans to the poor across low- and middle-income countries.

I have written essays describing how researchers and policymakers can respond to emerging technological shifts. **AI is Transforming the Economy; Understanding its Impact Requires both Data and Imagination** (*Nature*, 2025) describes two challenges: experiments on AI tend to be out of date by the time they are published, and experiments hold factors fixed that will change in the real world. I outline ways that social scientists can contribute to the study of AI more quickly: social science fiction (as termed by Jean Tirole), forward looking experiments, and piloting redesigned portions of the economy to provide advance notice of how they will reequilibrate. Other writing has laid out the terrain of the implications of AI for low-income countries. **Artificial Intelligence for the Poor** (*Foreign Affairs* website, 2023) was featured in a lead story in *The Economist* and was featured by the *BBC World Service*. I also wrote **Technology for Development** (IMF Finance & Development, 2023). Additionally, **To Regulate Network-Based Platforms, Look at Their Data** (with Chiara Farronato, *Harvard Business Review* website, 2021) describes how digital regulators can evaluate emerging antitrust questions using data that firms collect automatically as a side effect of operation.

I have worked on policy to ensure that benefits of AI are accessible to low-income populations around the world. I was one of 4 academics from around the world asked to brief 20 Central Bank Governors on the economic impacts of AI at the Bank of International Settlements in Basel, Switzerland in 2024. The US State Department circulated my writings, and invited me to two UNGA side events with Secretary of State Antony Blinken, the second of which unveiled a $100m initiative on AI for sustainable development with executives of AI firms. Based on my work, I was asked to join the Center for Global Development, where I now serve as a Nonresident Fellow. My policy work includes developing policy to reduce the barrier of communication costs for access to AI, and proposing the design of data markets to cover underrepresented languages. I regularly speak at the World Bank and have been asked to contribute to multiple World Development Reports (as a panelist for 2021 and academic advisor for 2026) and to join the advisory committee for the flagship report on the digital economy for Africa. I have been invited to present at the US Federal Trade Commission (FTC) and Federal Communications Commission (FCC), and my papers on network effects have been cited extensively in private legal disputes over network firms.

With Nick Tsivanidis I created a multiparty collaboration with a large telecom in Nigeria to produce mobile phone mobility estimates for the Lagos Metropolitan Area Transport Authority. Those estimates are being used to guide the city's transit plan.

My research has had impact on the research funding ecosystem. The expansion of digital credit has spurred the Gates Foundation to fund several multimillion-dollar initiatives (the Digital Credit Observatory at UC Berkeley, and the Consumer Protection Initiative at IPA). Due to my work on the supply of digital services, I was asked to serve on several related multimillion-dollar Gates Foundation initiatives, including *Financial Inclusion Through Interoperability* (Toulouse School of Economics, affiliate) and *Retail Finance Distribution* (University of Ghana, scientific committee).

From 2016 to 2020 I served on the board of the nonprofit I cofounded with Michael Kremer and Shawn Cole, Precision Agriculture for Development (now Precision Development), which currently serves over 4 million farmers with digital advisory services.

Teaching

A popular view suggests that theory is becoming less relevant: with modern datasets we can just 'let the data speak.' I developed and taught the master's level course **Applying Machine Learning** (DSPC IA7100), which begins with this premise and quickly reveals that it is difficult to learn from data without imposing any structure. The arc of the course then follows the question: 'what do you need to add to data in order to learn from it?' I created a series of assignments that mimic real world analyses to explore this, contrasting machine learning approaches (trees, random forests) with more parametric approaches (OLS, LASSO), addressing overfitting, uncovering the conditions under which we can interpret estimates as causal, and finally exploring strategic interaction with structural models. We explore each topic through datasets designed to mimic real decisions made by policymakers and managers. A highlight of the course is a week-long pricing competition in which students set daily prices to compete with another group, learning from historic data, economic models, and from each others' behavior. I currently teach the course to Master's students, but I originally developed it for undergraduates, and taught it at Brown University for many years. The course was one of the first combining machine learning and economics at any institution, and has influenced the design of other courses.

I developed and currently teach the master's level course **Our AI Future** (DSPC IA7175). Our institutions were developed in a context with different technologies: where travel and communication were slow and expensive, and thinking had to be done by humans. New technologies afford—and may require—different ways of organizing society. The course considers historical episodes of technological change and our current era, following how shifts in technology can shift the economy and society. I use the course itself as a laboratory to explore the impacts of AI on education. The course then considers how AI may reshape other sectors, including governance and transportation; and the cross-cutting questions it raises about values, economic wellbeing, and purpose.

In addition to my current courses, I have taught a course on **Applied Causal Inference** for undergraduates at Brown University (Econ 1629), which uses inquiry-based assignments to illustrate core principles. I also developed and taught a semester PhD course **Development Economics** at Brown University (Econ 2520) which exposes students to traditional topics using a sequence of intuitive models, and also includes unique sessions on passively collected big data and creating research partnerships. These sessions have shaped the research of several graduate students.

Service

At Columbia SIPA, I have guided the development of the new Data Science for Policy concentration. I developed the two SIPA classes mentioned above to meet the need for deep expertise at the intersection of policy and machine learning/AI, and guided the creation of a prerequisite course, *Advanced Computing for Policy*. I have contributed regularly to the school's discussion around AI, by providing input into the school's AI policy, engaging with students and fellows at the Institute of Global Politics, and serving on panels at events for admitted students, alumni, and the faculty.

I regularly mentor undergraduate, master's, and PhD students across SIPA and economics, and served on the committee for 10 PhD's between Columbia and Brown University. I regularly include students as coauthors, and have coauthored with 9 PhD students, 1 master's student, and 2 undergraduates. At Columbia, in 2025 I started a reading group on Artificial Intelligence and Economics with Augustin Chaintreau, with students across economics and computer science. In 2024-5 I co-organized the Development Colloquium for the SIPA Center for Development Economics and Policy with PhD students from SIPA and across Columbia. I regularly attend and contribute to the Data Science Initiative's Computational Social Science Working Group.

Part of my service to the profession includes organizing. I co-organized the *Bravo Center Workshop on Economics of Algorithms* in 2022, the *NBER Digital Economics and AI Meeting* in Spring 2025, and the session, *Economic Applications of Machine Learning*, at the American Economic Association 2018 meetings. Despite disrupted attendance at the conference due to a blizzard, the session was standing room only; people had to be turned away. I also served on the committee for various computer science conferences (track chair once, program committee 4 times, reviewer twice), and a reviewer for top journals and funders, as listed on my CV.

References

My papers (with coauthors) cited in this statement

Amaya, S.R., Björkegren, D., J. Blumenstock, S. Nair, 2023. Credit Scores That Prioritize Customer Welfare: Theory and Evidence from Nigeria. International Conference on Learning Representations (ICLR 2023) Workshop on Practical Machine Learning for Developing Countries, Kigali, Rwanda.

Bernheim, B.D., Björkegren, D., Naecker, J., Pollmann, M., 2021. Causal Inference from Hypothetical Evaluations (Working Paper No. 29616), Working Paper Series. National Bureau of Economic Research.

Björkegren, D., 2022. Competition in Network Industries: Evidence from the Rwandan Mobile Phone Network. The RAND Journal of Economics 53, 200–225.

Björkegren, D., 2026. Market Beliefs about Open vs. Closed AI. Working Paper.

Björkegren, D., 2020. Measuring Informal Work with Digital Traces: Mobile Payphone Operators in Rwanda, in: Proceedings of the 2020 International Conference on Information and Communication Technologies and Development, ICTD2020. Association for Computing Machinery, pp. 1–5.

Björkegren, D., 2019. The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda. Review of Economic Studies 86, 1033–1060.

Björkegren, D., Blumenstock, J.E., Folajimi-Senjobi, O., Mauro, J., Nair, S.R., 2025. Welfare Effects of Digital Credit: A Randomized Evaluation in Nigeria. Economic Development and Cultural Change 74, 117–140.

Björkegren, D., Blumenstock, J.E., Knight, S., 2025. What do Policies Value? *Forthcoming*, Review of Economic Studies.

Björkegren, D., Blumenstock, J.E., Knight, S., 2026. Manipulation-Robust Prediction. American Economic Review, *conditionally accepted (data)*.

Björkegren, D., Choi, J.H., Budihal, D., Sobhani, D., Garrod, O., Atherton, P., 2025. Could AI Leapfrog the Web? Evidence from Teachers in Sierra Leone. Working Paper.

Björkegren, D., Duhaut, A., Nagpal, G., Tsivanidis, N., 2024. Public and Private Transit: Evidence from Lagos. Working Paper.

Björkegren, D., Grissen, D., 2019. Behavior Revealed in Mobile Phone Usage Predicts Credit Repayment. World Bank Economic Review.

Björkegren, D., Karaca, B.C., 2022. Network Adoption Subsidies: A Digital Evaluation of a Rural Mobile Phone Program in Rwanda. Journal of Development Economics 154, 102762.

Choi, J.H., Garrod, O., Atherton, P., Joyce-Gibbons, A., Mason-Sesay, M., Björkegren, D., 2024. Are LLMs Useful in the Poorest Schools? TheTeacher.AI in Sierra Leone, in: Proceedings of the 2023 Neural Information Processing Systems (NeurIPS) Workshop on Generative AI for Education.

Milusheva, S., Bjorkegren, D., Viotti, L., 2021. Assessing Bias in Smartphone Mobility Estimates in Low Income Countries, in: ACM SIGCAS Conference on Computing and Sustainable Societies, COMPASS '21. Association for Computing Machinery, New York, NY, USA, pp. 364–378.

Rolf, E., Simchowitz, M., Dean, S., Liu, L.T., Björkegren, D., Hardt, M., Blumenstock, J., 2020. Balancing Competing Objectives with Noisy Data: Score-Based Classifiers for Welfare-Aware Machine Learning. Presented at the International Conference on Machine Learning (ICML).

Not peer reviewed

Björkegren, D., 2010. "Big data" for development.
Björkegren, D., 2023. Artificial Intelligence for the Poor. Foreign Affairs.
Björkegren, D., 2025. AI is Transforming the Economy — Understanding its Impact Requires both Data and Imagination. Nature 648, 535–537.
Björkegren, D., Blumenstock, J., 2023. Technology for Development. IMF Finance & Development.
Björkegren, D., Duhaut, A., Mbonu, O., Nagpal, G., Tsivanidis, N., 2026. Representation in Mobility Data in Emerging Cities. American Economic Association Papers and Proceedings.
Björkegren, D., Farronato, C., 2021. To Regulate Network-Based Platforms, Look at Their Data. Harvard Business Review. Website.
Björkegren, D., Grissen, D., 2018. The Potential of Digital Credit to Bank the Poor. American Economic Association Papers and Proceedings.

Other references

Andrews, I., Farboodi, M., 2025. Do Markets Believe in Transformative AI? Working Paper Series. https://doi.org/10.3386/w34243
Athey, S., 2018. The Impact of Machine Learning on Economics. The Economics of Artificial Intelligence: An Agenda 507–547.
Bharadwaj, P., Jack, W., Suri, T., 2019. Fintech and Household Resilience to Shocks: Evidence from Digital Loans in Kenya (Working Paper No. 25604). National Bureau of Economic Research. https://doi.org/10.3386/w25604
Brailovskaya, V., Dupas, P., Robinson, J., Robinson, J., 2021. Digital Credit: Filling a hole, or digging a hole? Evidence from Malawi (Working Paper).
Bruckner, M., Scheffer, T., 2011. Stackelberg Games for Adversarial Prediction Problems, in: Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '11. ACM, New York, NY, USA, pp. 547–555. https://doi.org/10.1145/2020408.2020495
Chow, T., Halperin, B., Mazlish, J.Z., 2025. Transformative AI, existential risk, and real interest rates. Working Paper.
Donovan, K.P., Park, E., 2019. Perpetual Debt in the Silicon Savannah [WWW Document]. Boston Review. URL https://bostonreview.net/class-inequality-global-justice/kevin-p-donovan-emma-park-perpetual-debt-silicon-savannah (accessed 8.4.21).
GSMA, 2008. Taxation and the Growth of Mobile Services in Sub-Saharan Africa.
Hardt, M., Megiddo, N., Papadimitriou, C., Wootters, M., 2016. Strategic Classification, in: Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, ITCS '16. ACM, New York, NY, USA, pp. 111–122. https://doi.org/10.1145/2840728.2840730
Lucas, R.E., 1976. Econometric policy evaluation: A critique. Carnegie-Rochester Conference Series on Public Policy 1, 19–46. https://doi.org/10.1016/S0167-2231(76)80003-6
Zuboff, S., 2019. The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power, 1st edition. ed. PublicAffairs.