

Nicholas Emery-Xu
University of California, Los Angeles
October 2026

Research Statement

My research seeks to understand the benefits and costs of the development and diffusion of information technologies and to analyze relevant policies to increase welfare. To answer my research questions, I combine novel data with contemporary structural modeling techniques. Where data is lacking, I rely on economic theory, guided by careful analysis of relevant forces in case studies. My first field, situated in empirical industrial organization, considers the impacts of digital technology adoption on firm productivity and barriers affecting the rate of diffusion of such technologies. Because many of these technologies are dual use, my second field, situated in political economy, seeks to understand how such technologies impact international conflict and cooperation.

I. Digital technology adoption and firm production

My primary research field is concerned with measuring the impacts of digital technology adoption on aspects of firm productivity. Such technologies are general-purpose, having the potential to impact firm growth and dynamism across many sectors of the economy. In addition, they often diffuse slowly due to a combination of supplier market power, adoption frictions, and learning-by-doing effects.

In my job market paper, “Firm in the clouds,” I study the effects of supplier market power on downstream firm cloud computing (infrastructure-as-a-service) adoption, productivity, and growth. I develop a model of downstream industry dynamics, in which firms adopt computing from any of several oligopolistic cloud providers, who compete in Bertrand competition to set prices, or choose to retain on-premise computing. To model a realistic account of adoption and supplier choice, I extend prior models of firm technology adoption and model product heterogeneity as differentially affecting firm productivity and maintenance costs. I allow products to differ by quality and complementarity with firms’ on-premise server software, and price. This paper combines French data from Orbis on firm inputs and output with data on technology provider adoption from Aberdeen, a market research firm. To this data I add a novel panel dataset of annual datacenter level average prices and CPU quality gathered from cloud providers’ websites, Microsoft Azure virtual machine trace data, and third party websites.

I estimate the production function using a control function approach with non-IT investment as my proxy variable. I find that cloud computing adoption increases firm productivity by a small but statistically significant amount – 0.2% to 1.6% - with larger gains in manufacturing than in services and increased gains for firms that multihome. Using these estimates, I simulate a model of industry dynamics in which downstream firms choose bundles of computing providers as well as their capital stock. Cloud computing providers compete on prices. Because of the large number of firms adopting the technology, I employ the Markov Moment Equilibrium developed by Ifrach and Weintraub (2017), permitting downstream firms to keep track of their own state

and input prices. Simulating the model, I quantify how adoption is being driven by switching frictions, falling entry costs, and a high degree of complementarity among Microsoft products. Using my estimates, I first compute how much cloud provider market power inhibits the profit of downstream firms. Studying the direct impacts of market power, I find that a hypothetical merger between Microsoft's and Google's cloud divisions would reduce total surplus by €905 M annual (0.29%). The EU Data Act's egress fee ban increases welfare by €193 M (0.06%) per year. Increasing the complementarity of Microsoft server software with rival cloud platforms increases price competition, increasing welfare by 0.16%.

The next major general-purpose technology that promises to revolutionize firm production is deep learning, the technology underlying modern artificial intelligence systems. In "Economic impacts of AI-augmented R&D," we use a simple model of endogenous productivity through innovation to show that, if physical capital accumulation for R&D is more rapid than human capital accumulation, then a shift to a more highly-capital intensive production process for R&D should increase the productivity growth rate of firms. Using the computer science literature on scaling laws for neural networks, in which model performance scales with pretraining compute, we argue that deep learning may accelerate the innovation process and thus productivity growth.

We test this model using data on deep learning models for two computer vision tasks. We augment the dataset on compute inputs from Thompson et al. (2020) and create a novel measure of human capital by training a neural network to predict the forward citations and journal ranking of all authors on a given computer science paper, using information such as authors' departmental grants, institution, and prior publication record. We train this model on a training set of all published computer science authors on arXiv and Scopus, showing that it outperforms more traditional statistical learning techniques such as LASSO on our hold-out set.

We apply this model to the deep learning models in our dataset. Then, we estimate a production function for deep learning, using pre-training compute and our human capital measure as inputs and model error rate as our output. We find that deep learning has a compute capital cost share of 35-40%, higher than the STEM R&D average of 7%. Applying our results to a calibrated model of the US economy, we find that if deep learning-augmented R&D were to fully diffuse, the productivity growth rate would rise from an average of 1.2% annually to between 2.1% and 2.4%.

Moving from hypothetical to measured AI adoption, in "Artificial intelligence, productivity, and the labor share," I study the impact of deep learning adoption among Spanish manufacturing firms. I ask how deep learning has diffused among Spanish manufacturing firms, finding a slow but steady rise from 6% of firms in 2018 to 12% by 2022, with higher adoption rates among automobile, computer, and other high-technology manufacturing sectors. Data comes from the ESEE.

I build and estimate a model of firm production with labor-augmenting and Hicks-neutral productivity. I find that AI adoption increases labor productivity by 15%-17%, about thrice the size of the effect of robot adoption, and that there is no statistically significant effect on Hicks-neutral productivity. Using my estimates, I plan to simulate the impact of AI on productivity and the labor share over the next decade.

Several ongoing projects will continue this line of research. First, a set of projects with Neil Thompson seeks to understand the role of computing progress on the productivity boom in the United States in the 1990s and 2000s and subsequent decline. The first hypothesizes that the slowdown may either be due to a slowdown in computing progress or to non-IT bottlenecks in production. We build a model of production with IT-augmenting productivity. We plan to use data from the Annual Survey of Manufacturers and Census of Manufacturers to estimate the model. The second project uses CPU usage data from over 1 million computers running on Intel chips. We use the data to study what fraction of improved CPU performance improves user performance and what fraction is misallocated. In process analysis is studying results for the graphics and software development industries.

II. Technology, arms racing, and war

In my secondary research field, I study the impact of dual use technology development and adoption on arms racing, conflict, and cooperation.

In Emery-Xu, Park, and Trager (2023), we study the impact of information on welfare in a model of risky technology development. In many national security contexts, there is a tradeoff between investing in technological progress to develop ahead of rivals and investing in the safety of that technology. We show that these forces are relevant to both the development of the atomic bomb and current policy debates over the development of advanced AI models.

Motivated by these cases, we develop a model of an R&D race for a risky technology, in which devoting effort to safety lowers the chance of a disaster affecting all agents but lowers the change that a given agent will affect the race. We compare total welfare under two main cases: 1) private information about agent resource endowment and 2) common knowledge about such endowments. We find that the relative welfare of the two cases depends on the decisiveness of the race. When the contest is noisy, private information is safer; otherwise, common knowledge is safer.

In Bressler, Emery-Xu, Trager, and Dafoe, we study the impact of shifts in the offense-defense balance on the welfare of two agents at war. We find that welfare is U-shaped under offensive advantage. When technology is defense-dominant, arms expenditures are low and welfare is high because attacking is costly. As the marginal returns to arming for attack or defense approach parity, peace requires higher arms levels to deter an attack, while conflicts become longer. In either case, welfare falls. Finally, if technology becomes highly offense-dominant, the only sustainable equilibrium is a war in which the attacker quickly wins in expectation. Case studies from feudal Japan and imperial China confirm that military innovations can produce such offense-dominant welfare increases.

Publications and working papers by topic

Here I list my working papers and publications by topic:

I. Firm IT adoption and productivity

1. Emery-Xu, N. (2026). "Firm in the clouds: Supplier market power and technology adoption in cloud computing." *Job market paper*.
 2. Besiroglu, T., N. Emery-Xu, and N. Thompson. (2024). "Economic impacts of AI-augmented R&D." *Research Policy*. 3(7):105037.
 3. Emery-Xu, N. (2026). "Artificial intelligence, productivity, and the labor share." *Working paper*.
- II. Technology, arms racing, and war
1. Emery-Xu, N., A. Park, and R. F. Trager. (2024). "Uncertainty, information, and risk in international technology races." *Journal of Conflict Resolution*, 68(10):2019-2047.
 2. Bressler, R. D., N. Emery-Xu, R. F. Trager, and Allan Dafoe. (2026). "The offense-defense balance and the costs of anarchy: Welfare is U-shaped under offensive advantage." *Working paper*.

Sources

1. Ifrach, B. and G. Y. Weintraub. (2017). "A framework for dynamic oligopoly in concentrated industries." *The Review of Economic Studies*, 84(3):1106-1150.