Research Statement

Stan de Ruijter
University College London

November 2026

Technological change, trade, and other economic shocks can affect workers even when they are not directly exposed to them. A supplier may redesign a product so that workers at its buyers perform fewer tasks, or entirely new ones. A trade shock can likewise reach industries and places with little direct exposure to foreign trade competition. Modern economies are made up of dense production networks, but economists have mostly focused their research efforts on those workers, firms, and regions which experience these shocks directly. This narrowed focus has obscured the broader and more diffuse effects of economic shocks.

My research asks what shocks do in networked economies and how economists should study them. I study not just the direct effects of these economic shocks but their indirect effects as well: how they travel through the relationships that organize production in modern economies. This agenda lies at the intersection of labor economics, international trade, and production networks. I combine administrative data, empirical research designs, and economic models to study how changes in technology and trade reshape the demand for labor and the distribution of economic opportunities. Three papers form the foundation of this program: on technology transmitted through supply chains, on shift-share designs under network propagation, and on the indirect effects of the China shock. Together, my research shows that approaches which focus on the direct effects of economic shocks miss the substantial indirect effects on network-connected workers, firms, and industries, and misstate who bears the costs and who gains.

Technology in production networks: indirect effects on workers and firms

My job market paper, "Displacing and Reinstating Technologies in Production Networks: Workers and Firms," asks how innovation at a supplier affects the firms and workers that buy from it. Consider a car manufacturer that replaces serviceable components with sealed, pre-assembled modules. Workers at repair shops which buy those components may no longer need to diagnose or repair the underlying machinery, reducing demand for their labor. By contrast, innovations such as electric drivetrains or sensor systems can create downstream tasks in calibration, maintenance, and software. In both cases, the repair shop's workers are affected without their employer itself adopting the technology. This paper extends the task-based approach to technological change across firm boundaries: innovations embodied in purchased inputs can either displace tasks at buying firms or reinstate labor in new tasks. I show that the type of technology decides the direction of the effect. When suppliers develop displacing technology, buying firms shrink and their labor share falls, and workers who started in the exposed industries earn less wherever they work later; under reinstating technology, each of these effects is reversed.

To study these effects, I use Danish matched employer–employee registers covering the universe of Danish workers from 1980 to 2023. I link firms and workers to measures of technological change and granular input–output tables. Multiple independent patent classifications distinguish labor-displacing from labor-reinstating innovations. I construct each industry's exposure to the innovations of its supplier industries, using predetermined input shares. For workers, I assign exposure using their first observed industry and follow their earnings across subsequent jobs. This allows me to study the consequences of the industries in which workers acquired their experience rather than conflating the effects of shocks with their decisions to move. The main empirical challenge is distinguishing supplier innovation from broader trends in workers' original industries. I show that identifying effects on buying industries requires shocks to supplying industries to be as good as random, in the sense of Borusyak, Hull, and Jaravel (2022). Because the patent stocks are built from US patents, in which Danish inventors account for a negligible share, and input shares are fixed before the period of study, the shocks are unlikely to reflect conditions in Danish buying industries. On correlated trends, tests using future exposure and pre-period industry trends are broadly reassuring.

Buying firms exposed to suppliers' displacing technology reduce employment and labor shares, while value added per worker and operating margins rise. Adjustment occurs primarily through reduced hiring; wages paid to incumbent workers and new hires at those firms are unchanged. Reinstating innovations produce effects in the opposite direction: firms downstream of them see increases in employment and labor shares. Suppliers' technology is thus consequential for downstream labor demand, even without a corresponding rise in the buyer's measured capital stock. Its effect on workers' pay is economically comparable in magnitude to that of technological change in their own industry.

For workers, the consequences of upstream technological change follow them well beyond their original employers: those whose initial industries are exposed to suppliers' displacing innovations subsequently earn less, including after they have moved to other industries. The estimated earnings loss is as large among workers who left their original industry as for those who stayed. Part of the adjustment occurs through moves to lower-paying employers, which account for about two-fifths of the wage effect, alongside lower employment and more exits from the labor force. Losses in hourly pay are concentrated among relatively well-paid white-collar workers, while lower-paid workers adjust more through reductions in hours and employment. These patterns suggest a possible 'vintage penalty': skills developed around an earlier organization of production may lose value as suppliers change what downstream firms need workers to do. Meanwhile, workers indirectly exposed to reinstating technology see higher pay throughout their careers.

To interpret the firm responses, I develop a task-based model in which inputs and workers perform different parts of downstream production. A shift of suppliers' innovation toward displacing technologies has two effects on the buyer at once: the inputs perform more of the buyer's tasks, and more cheaply, and part of what the buyer used to sell is now produced by others. The buying firms' accounts discipline the model: in response to a shift toward upstream displacing innovation, purchases contract most, margins rise, and output per buyer falls while the number of buyers is unchanged. Fitting the model's five parameters to fourteen moments of the buying firms' accounts, their wage, and their output price, I find that, per unit of the suppliers' shift toward displacing (process) innovation, the buying industry's demand falls by 0.79 (s.e. 0.10) and the input's cost per task by 0.17 (s.e. 0.23). The demand shift accounts for the entirety of the fall in output per firm, and for at least half of the fall in employment and pay. Taken together, the results suggest that the labor-market costs of technological change can arise outside adopting firms and persist across workers' careers. For adjustment policy, the losses from displacing technology arrive one step down the supply chain, as fewer hires and lower-paying later jobs for workers who started in the exposed industries, not as layoffs at the firms that use the technology, which is where support for displaced workers is directed.

Shift-share designs under network propagation

My second paper, "Shift-Share Designs with Network Propagation," studies a related problem in applied econometrics: how should researchers measure a region's exposure to an industry shock when the shock is propagated through a network? Conventional shift-share, or Bartik, measures weight industry shocks by local employment shares. When production linkages transmit those shocks to other industries, a place's exposure will depend on more than its employment in the directly affected sectors. The paper characterizes the resulting omitted network-exposure term, derives a closed-form diagnostic computable in advance from network weights and employment shares, and develops a correction which nests the standard estimator.

Applied to the canonical China-shock design of Autor, Dorn, and Hanson, the standard 2SLS coefficient is 2.4 times the corrected one, yet the two deliver nearly identical aggregate predicted impacts, a cancellation that holds under an exact and falsifiable condition the paper characterizes. But the geographic incidence of predicted impacts is substantially different. Thirty-four of the 100 most exposed commuting zones by the conventional measure, and 68% of the top 25, drop out of the corresponding corrected rankings, as network exposure shifts estimated impacts toward industrial regions linked to directly affected sectors through their supply chains. On held-out geographies, the corrected exposure triples out-of-sample predictive fit, with an R^2 of 0.26 against 0.08. Supply-chain-exposed places, like the steel, chemicals, and petroleum regions of the industrial Midwest, carry impacts that the standard measure attributes to the directly affected textile belt. Identifying the places affected by a shock is therefore not the same as identifying the places most directly exposed, which matters for place-based adjustment policies.

The China shock through the production network

In "Decomposing the China Shock: Direct and Upstream Effects in the United States and Germany," joint work with Christian Dustmann, we ask how China's rise affected two major advanced economies through their domestic and international production relationships. Beyond direct import competition, an industry can be affected when its customers expand or contract, including in other countries, and accounting for these connections may help explain why the United States and Germany experienced contrasting consequences of China's integration into world trade. We use a computable general equilibrium model in the Caliendo–Parro tradition to compare observed trade integration with a counterfactual in which trade with China remains at its 2000 level. We introduce novel decompositions of the model-implied changes in production and labor demand along three dimensions: direct trade effects against effects induced upstream, in the industries that supply the directly affected producers; final against intermediate goods; and domestic against international channels. We find that most of the impact arrives through upstream linkages rather than the direct flows. In the United States, the upstream effects more than offset the losses from direct import competition. In Germany, where the direct effect of Chinese trade was already positive, the upstream effects are almost twice as large again. Third-country linkages play a substantial role: for example, for German machine tools, more than two-thirds of the upstream production effect originates in other countries' trade with China. The implications also differ across workers: male manufacturing workers lose disproportionately in the United States but gain in Germany, while service workers benefit from upstream linkages in both. The project shows how countries' positions in global production networks shape the incidence of a common trade shock, beyond what bilateral trade exposure alone would suggest.

Ongoing and future work

My current work has led me toward a broader question about labor-market adjustment: why do shocks continue to affect workers after they leave the industries in which they were initially exposed? The persistence of wage losses in my job market paper points toward human-capital specificity, but also toward changes in sorting across employers. I plan to use Danish administrative data to investigate these mechanisms through workers' occupational histories, transitions between firms with different wage premia, and the changing returns to experience in exposed industries. A central objective is to distinguish whether workers lose the value of skills acquired under older production methods or instead face persistent barriers to finding similarly productive employers, which connects network-mediated technological change to debates about skill transferability and the long-run costs of displacement.

A related line of work will examine the interaction between technological adoption and indirect exposure. My existing paper measures innovations originating in supplying industries; I want to bring firm-level adoption measures alongside these industry measures and study when the two reinforce or offset each other. I am also interested in worker flows between adopting and non-adopting firms as a possible transmission channel. These projects would help establish when the technologies embodied in inputs, direct adoption by employers, and the movement of workers are complements or substitutes in spreading the effects of innovation. They would also speak to why some of the largest losses appear among highly paid workers in occupations that coordinate production, rather than in the routine jobs often emphasized in discussions of automation.

Finally, I aim to develop the measurement agenda beyond the China shock. Energy-price changes and industrial policies are settings in which direct industry exposure may be a poor guide to the incidence of a shock. In "The Network Incidence of the 2022 European Energy Crisis," a first step, I find that about half of the cost shock from the crisis reached industries through their suppliers rather than through their own energy use. I want to establish when network-adjusted measures materially alter empirical estimates, when conventional measures remain adequate, and whether the differences help predict subsequent outcomes for workers and places. A longer-term application is the targeting of adjustment assistance: if workers and communities can be indirectly exposed, how much does a direct-exposure rule miss, and what would a better-targeted rule require? Across these projects, my goal is to understand the mechanisms that transmit economic change, and to develop evidence and tools that identify its consequences for workers and places beyond those who are directly exposed.