

2018

Essays on the Economics of Labor Demand and Policy Incidence

Andrew Garin
Harvard University

Essays on the Economics of Labor Demand and Policy Incidence

Andrew Garin

In the United States, a common policy strategy to expand employment and increase wages is to try to expand the demand for labor, either through support to employers in the form of subsidies or tax breaks or through direct government purchases. However, while demand-side policies should increase wages and/or employment in theory, it is often unclear who benefits from such policies in practice. In competitive labor markets, the benefits of employment subsidies may accrue largely to employers or to other factors of production besides the target worker population. It is thus crucial to empirically assess the real-world effects of such demand-side interventions using evidence from past experience, and, moreover, to identify the economic principles that determine why different policies may succeed or fail at achieving their desired impacts. This dissertation comprises three independent essays that study demand-side interventions and shocks to the labor market and evaluate who wins and who loses in practice. I draw on novel data to identify natural experiments that generated as-good-as-random changes in labor demand conditions across both firms and regional labor markets. In each setting, I study who bears the benefit and burden of a different shock to the labor market, and, using economic theory as a guide, I highlight the principles that determine when demand-side policies successfully benefit the workforce.

In the first essay (joint with Filipe Silverio), I study how firms mediate the benefits of labor-demand expansions that come from product-market-side incentives. When an employer experiences an increase in product demand or a subsidy to produce more—which in turn increases its demand for labor—does it raise wages for current employees? Or, are wages pinned down by competition from other employers in that labor market? To answer this question, we study unique data on Portuguese exporters and their employees, in which we can identify cases where individual firms did comparatively better or worse than their similar competitors, based on where their customers were during the global Great Recession in 2008 and 2009.

In the second and third essays I study the effects of two significant public infrastructure investment efforts on regional labor markets in the United States, shedding light on when such policies can bolster wages and job creation in target regions. One essay studies the short-run effectiveness of stimulus-funded highway construction projects as a means of boosting local employment—construction-sector employment in particular—in stagnating regions during the Great Recession. The other takes a long-run view and examines how the rapid construction of war plants during World War

II in unusual locations (which were later reconverted for private civilian product and operated for decades afterwards) impacted regional manufacturing development, and, in turn, the availability of high-paying manufacturing jobs during the postwar era. In each of these studies, I collect new, geographically detailed, investment-level data sets that were crucial to conduct causal evaluations of each intervention.

Chapter 1

Do Wage Increases Reflect Firm-Level Labor Demand or Market Competition? Evidence from Idiosyncratic Export Demand Shocks

(with Filipe Silverio)

What induces employers to raise their employees' wages: their internal need for labor or competition from other employers in the labor market? In the canonical competitive labor market model, wages are determined by competition among employees and employers in labor markets, and firms always pay the same amount to identical workers. However, in real-world labor markets, where replacing employees and finding new jobs can be costly, workers at firms facing expansionary demand on the product market side may be able to negotiate higher wages, regardless of the state of the outside labor market. Understanding how and why employers behave in practice is important for assessing the impact of labor market interventions that target specific employers in large, global labor markets. Moreover, if wages are determined noncompetitively within firms in separate "internal labor markets" (Doeringer and Piore 1971), this might also in turn account for the well-documented role of firms in pay inequality found around the globe (Abowd, Kramarz, and Margolis 1999; Barth et al. 2016; Card et al. 2016; Song et al. 2015).

In this paper, we study how employers respond to changes in product market conditions that incentivize them to produce. Our goal is to test how much and under what circumstances firms raise wages to match heightened production goals—even when wage competition from rival employers remains unchanged. Such a test is difficult to implement in practice, as one needs to isolate a source of variation in demand conditions that affects output, employment, and wages but that is also, first, uncorrelated with any changes in the skills and effort employees bring to the table, and second, not associated with a general change in demand among rival employers. Several studies to date examine how changes in wage correlate with changes in output or productivity (Alvarez et al. 2018; Card, Devicienti, and Maida 2013; Guiso, Pistaferri, and Schivardi 2005); however, the conclusions that can be drawn from such studies are limited by the possibility that changes in output are driven either by labor supply (rather than labor demand) or by changes in labor

demand that are common to entire markets. This paper attempts to fill that gap.

To assess the extent to which wages respond to firms' internal labor demand rather than labor market competition, this paper studies a novel natural experiment that idiosyncratically changed the individual firms' product demand—without affecting conditions at their closest labor-market competitors. We study how much Portuguese exporters' ability to sell during the global Great Recession was impacted by where—but not what—they had been selling before 2008. As different countries' import demands for any given product were jolted in unpredictable ways during this turbulent period, otherwise similar exporters experienced unexpectedly differential changes in product demand based solely on who their customers were—even those selling the same product.

We follow a budding strain in the literature on international trade (Berman, Berthou, and Héricourt 2015; Hummels et al. 2014; Mayer, Melitz, and Ottaviano 2016) and use very detailed data on how much of each detailed product variety each firm exported to each overseas destination to create measures of exposure to each foreign market. We then combine these exposure measures with observed changes in importing at each of those destinations during the Great Recession to create a measure of how much each firm's export demand changed based on overseas conditions. Importantly, we implement a novel decomposition of the export demand change into two components: one reflecting global changes in product demand that might have affected many firms in a product market, and another that measures how much purchases of that product in the specific destination changed compared to other countries.

We show that while both components have meaningful impacts on firms' product and labor demand conditions, as evidenced by sales and payrolls, the former “common” component also affects the sales and payrolls of other close competitors. However, the second “idiosyncratic” component of an exporter's demand shock impacts that specific firm, but it otherwise has no effect on close competitors. This offers precisely the sort of quasi-experimental differences in demand across firms necessary to test how much internal demand incentivizes firms to adjust wages.

We find that, even in the absence of a marketwide demand shift, the wages of prerecession employees adjust significantly in response to firm-level demand changes. A demand shift resulting in 10 percent more output results in a 1.5 percent larger rise in hourly wages for those incumbent workers. These effects arise primarily within continuous employment spells at the initial employer, suggesting that these wage changes reflect bargaining power within the firm—and not an increase in general human capital.

In our Portuguese setting, this bargaining power may arise in part from institutional restrictions on firing. Consistent with high firing costs, we find that employers also respond to production demand shocks by hiring more or fewer work-

ers—but not by firing existing employees. In addition, we find that in sectors where wage effects exist, they are widely shared across workers of different pay level, skill levels, genders, and occupations.

However, wages do not respond in the same manner in all settings. Consistent with theories of firm-specific human capital (Becker 1962; Lazear 2009), costly searches for workers with unique skills (Mortensen and Pissarides 1994), and institutional firing costs (Lazear 1990), we find that wages only respond to firm-level demand in sectors with higher barriers to employee turnover, which we infer based on low quit rates and longer tenures. In sectors with high barriers to turnover we find that the wage incidence of firm demand shocks is quite high: A shock that changes output by 10 percent leads to a 3 percent wage increase. By contrast, in sectors with low turnover costs, wages are mostly pinned down by labor market competition, and we find no evidence that firms adjust wages based on their own demand conditions. These findings suggest that turnover costs play an important role in wage determination—the easier it is to replace workers, the less wages reflect internal labor demand within the firm as opposed to labor market competition.

These findings suggest that wages among otherwise identical workers can vary substantially due to circumstances at their firms that are beyond their control. Moreover, our estimates can reconcile a puzzle in the earlier literature of pay difference across firms: Card et al. (2016) find that while firms with 10 percent higher labor productivity levels have 1.5 percent higher pay premiums in cross-sectional comparisons, nonexperimental studies that correlated changes in pay with changes in firm output rarely find wage increases large enough to rationalize the general productivity-pay relationship. By contrast, in our experimental study, we find a causal relationship between firms' output growth and wages that is much larger than implied by correlation analyses (including our own), and that exactly matches the cross-sectional productivity-pay relationship. When we use our causal estimates to predict the amount of wage variance that should arise from differences in firm performance, we account for roughly the same share of the total variance in wages explained by firm-level pay premiums in the common decomposition from Abowd, Kramarz, and Margolis (1999). Thus, our estimates imply that changes in pay policy in response to employer performance can plausibly generate large cross-firm wage differentials—as long as there are sufficiently large barriers to replacing workers. These findings highlight the value of analyzing natural-experimental evidence in understanding the factors that drive wage determination.

Chapter 2

Putting America to Work, Where? Evidence on the Effectiveness of Infrastructure Construction as a Locally Targeted Employment Policy

This paper studies spending on public infrastructure construction projects—in particular, “shovel-ready” projects that can commence immediately upon funding—a frequently proposed policy intended to boost local construction employment and overall economic health. Such projects can target particular geographies in a highly visible way, making them intuitively appealing as way to boost employment in distressed areas. These projects inherently create a demand for construction laborers to go to work in precise locations. As a result, one might expect that approving a project in a place will increase local employment in the construction sector, and perhaps increase employment more broadly in turn as construction workers spend their earnings locally. Despite the intuitive appeal of these arguments, however, it is important to evaluate whether the data support them.

To test these hypotheses, I study the 2009 American Recovery and Reinvestment Act (the Recovery Act), which authorized \$27 billion for supplemental shovel-ready road construction projects that could commence promptly, with priority given to economically distressed areas. In contrast to standard federal road expenditures, the Recovery Act required detailed reporting about all stimulus road construction projects nationwide. This provides a unique opportunity to study the local employment effects of infrastructure spending.

Using the Recovery Act’s spatially detailed data on infrastructure spending, I test whether places that received relatively more funding experienced more favorable employment outcomes than those that received relatively less using a variable treatment intensity difference-in-differences design. To the extent that construction workers might have been engaged in other construction work in the absence of the stimulus measure, counting bodies at project sites is insufficient; one must determine causal effects of spending relative to the no-spending counterfactual. To this end, I use a rich local-level data set to consider the plausibility of a selection-on-observables methodology. Contrary to concerns that funds were systematically targeted to places with unobservable worse downturns, I find surprisingly little evidence of any targeting based on observable employment trends.

One advantage of studying government highway spending is that there is a clear transmission mechanism by which expenditures should affect local employment. The first-order “direct effect” should be on employment of construction workers involved in the projects themselves, and there should be a “local multiplier effect” on nonconstruction employment in the same region only to the extent that those

workers’ expenditures support additional jobs. For any local multiplier effect to be plausible, it is necessary to first establish a credible direct effect on the construction sector.

I find that highway construction did have a direct effect on construction employment at the county level. In particular, \$1.00 of additional Recovery Act spending on local construction increased local construction payrolls by \$0.30 during the five years after the act’s passage, almost exactly labor’s share of construction revenues nationwide. These labor market effects were largest in 2010 and dissipated gradually over the following years. I find no evidence of differential preperiod trends across differently treated counties, supporting the identification assumption that all effects are causal results of additional stimulus spending. The finding that the magnitude of the direct effect is roughly what one would expect with zero crowd-out suggests that targeted Recovery Act spending during the Great Recession did not crowd out other local construction. However, I find that commuting matters in that local spending only impacts local employment in more isolated locales with smaller populations and smaller fractions of residents who travel to outside counties for work. When I test for effects in nearby commuting origins or destinations, I find some evidence of commuting spillovers, but the estimates are highly imprecise.

When I test for general equilibrium effects on the total employment and payroll levels within locales, I find effects close to zero, with very wide confidence intervals across all specifications. I find evidence suggesting that places with less commuting penetration have larger total employment effects, but the results are too imprecise to conclude that there is a large local multiplier anywhere. In sum, although I do not find evidence of large local multipliers, I cannot rule them out. A local multiplier of zero could simply indicate that the multiplier effects of localized spending are highly spatially diffuse; yet, an imprecise zero estimate may also merely reflect a lack of statistical power. While the Recovery Act was a significant enough intervention to have a sizable effect on the construction sector in low-mobility counties, the local variation in highway spending may have been too small relative to baseline regional volatility to detect a local multiplier.

Importantly, although the magnitude of the effects on the construction sector is what one might expect in the absence of crowd-out, the implied cost per job I estimate is high relative to other estimates in the literature. My estimate of five construction job-years per \$1 million implies that one construction job-year requires \$200,000 in highway spending to be sustained. While I find a larger impact on construction employment compared to the state-level findings than in Leduc and Wilson (2017), the cost per construction job-year is much more expensive than the roughly \$30,000 cost per total job-year in Suarez Serrato and Wingender (2014) and Buchheim and Watzinger (2017). In order to reconcile the construction employment effects I find—which imply a cost

per construction job of \$200,000—with a \$30,000 cost per total job, a construction job would have to result in the creation of five additional nonconstruction jobs. While not ruled out by my estimated effects on total employment, this is not plausible based on prior studies of local employment multipliers, which typically find that an exogenous one job-year increase in low-skill work supports only one additional job in the same metropolitan area (Moretti 2010; van Dijk 2016).

Chapter 3

Public Investment and the Spread of “Good-Paying” Manufacturing Jobs: Evidence from World War II’s Big Plants

In this paper, I study the long-run local labor market effects of the publicly financed construction of large manufacturing facilities during World War II (WWII). Many have hypothesized that the postwar expansion of high-wage manufacturing employment and associated “Great Compression” in the earnings distribution was a direct result of the industrial mobilization for the war (Goldin and Margo 1992; Piketty 2014). This paper provides a direct test of this conjecture. In doing so, it speaks directly to a core question for economic policymakers: Can a brief public intervention have beneficial impacts on labor markets that persist well beyond the period of government intervention? Understanding the answer is vital to understanding whether industrial policy is justified in developing regions, whether place-based labor policies can help the middle class, and the extent to which infrastructure shapes how urban clusters form and persist. However, opportunities to study this question directly are extremely rare—big “pushes” are few and far between, and when they do occur in the form of plant openings or major infrastructure works, they are typically systematically targeted at places that are expected to grow or to stagnate.

The industrial expansion for WWII provides a unique opportunity to study this question. Due to the short-run military emergency, political and military leaders demanded that the United States increase its domestic industrial output nearly threefold over the course of only three or four years. This increment to output primarily consisted of airplanes, ships, ordnance (guns and ammunition of all varieties), explosives, and the metals and chemicals used in the production of those various types of matériel. Although the military attempted to incentivize firms to put their own capital on the line and build plants as necessary, for some particularly large plants in secure locations, these incentives were insufficient to attract any private investment, particularly when an expensive plant was built to churn out a product that was insufficiently similar to products more likely to be demanded during peacetime. In these cases, the plants were ordered and

owned by the U.S. government. As the private sector was unwilling to finance plant construction at the sites selected by the U.S. government, it is unlikely that similar plants would have been sited in the same locations if not for the war. These large, durable, public plants are more plausibly located for quasi-random short-run reasons than any comparable infrastructure investment in Western history.

To estimate the effects of siting a large plant in a specific locale, I compare counties that received large and completely federally funded plants to counties that were observably similar at the dawn of the Great Depression. My conjecture is that in the absence of a war, neither the control nor the treatment counties would have had such an additional plant open; the only reason potential outcomes differ across treatment and control counties were circumstances created by the war. I hone in on control groups using several methods, although the choice of method does not significantly impact the results. To estimate the effects of plant sitings I adopt a difference-in-differences approach: After verifying that labor market conditions in treatment and control locales evolved similarly prior to WWII, I study how much more employment and wages rose in the treated locales compared to the control locales. I find evidence that among similar counties, publicly financed plants were as good as randomly assigned. By contrast, plants that received substantial private funding seem to be systematically located in regions experiencing higher prewar industrial expansion. This supports my focus on government-financed plants.

I find that the construction of a large industrial plant during WWII had a large and persistent effect on the local labor market, with particularly pronounced effects on the manufacturing wage. After the postwar reconversion, manufacturing employment in counties grew by 30 percent and remained elevated until the 1980s. Moreover, the average wage of local production-line workers rose by 10 percent, and this resulting wage differential persisted beyond even the 1980s. I find that these manufacturing sector effects are associated with a general increase in median family incomes and, to a lesser extent, with higher wages in other sectors.

These findings highlight the large role that government-supported manufacturing expansions played in the spread of “good-paying” midcentury manufacturing jobs. These long-lived effects were the result of the initial investments during the war, not continued government or military investment. In fact, the most persistent effects were because of the groups of plants most likely to be reconverted for civil purposes—those making planes, tanks, and steel—rather than specialized ordnance plants that tended to cater to military production throughout the Cold War era. However, while government spending on investment in manufacturing capital complimentary to labor had large effects, I find that wartime government spending in general did not boost wages in a similar manner, consistent with earlier work by Fishback and Cullen (2013). The effects I identify are large in magnitude

relative to the costs: At a 10 percent local income tax rate, local governments would have recovered roughly half of the cost of the plant construction from the payroll increment alone.

This work adds to a growing literature evaluating the regional effects of place-based policies and large manufacturing investments. A study of “million dollar plants” that opened in the 1980s and 1990s (Greenstone, Hornbeck, and Moretti 2010) finds that while new plant sitings increase regional productivity, there were limited effects on local earnings; however, little is known about whether plant sitings had similar effects when manufacturing plants offered better work opportunities midcentury. Research on the Tennessee Valley Authority by Moretti and Kline (2014) finds that a major regional development push in the 1930s had long-run impacts, but it is difficult to conclude which aspects of the policy drove the results.

While the identification of persistent effects on average wages is of interest, the implications for policy depend on the reason for the wage increases. On one hand, the increase in local average wages may reflect an actual increase in earnings opportunities for certain individuals—those who would have resided near the plant in any circumstance and who might have gained access to better jobs and work experience with the plant siting. On the other hand, the rise in local wages might merely reflect an inflow of skilled workers drawn to the new plant that would have earned the same higher wage in any case, but perhaps in another location. In this latter scenario, the investment might only change where people work without affecting any single individual’s earning opportunities. I propose a method to distinguish between these two scenarios by studying the long-run effects of plant siting on the individuals residing in county before WWII, rather than annual county-level aggregates. I discuss plans to use longitudinally linked microdata to extend the study along these lines.

References

- Abowd, John A., Francis Kramarz, and David Margolis. 1999. “High Wage Workers and High Wage Firms.” *Econometrica* 67(2): 251–333.
- Alvarez, Jorge, Felipe Benguria, Niklas Engbom, and Christian Moser. 2018. “Firms and the Decline in Earnings Inequality in Brazil.” *American Economic Journal: Macroeconomics* 10(1): 1–43.
- Barth, Erling, Alex Bryson, James C. Davis, and Richard Freeman. 2016. “It’s Where You Work: Increases in Earnings Dispersion across Establishments and Individuals in the U.S.” *Journal of Labor Economics* 34(S2): S67–S97.
- Becker, Gary. 1962. “Investment in Human Capital: A Theoretical Analysis.” *Journal of Political Economy* 70(5): 9–49.
- Berman, Nicolas, Antoine Berthou, and Jerome Hericourt. 2015. “Export Dynamics and Sales at Home.” *Journal of International Economics* 96: 298–310.
- Buchheim, Lukas, and Martin Watzinger. 2017. “The Employment Effects of Countercyclical Infrastructure Investments.” CESifo Working Paper Series 6383. Munich: CESifo Group.
- Card, David, Anna Cardoso, Joerg Heining, and Patrick Kline. 2016. “Firms and Labor Market Inequality: Evidence and Some Theory.” Cambridge, MA: National Bureau of Economic Research.
- Card, David, Francesco Devicienti, and Agata Maida. 2013. “Rent-sharing, Holdup, and Wages: Evidence from Matched Panel Data.” *Review of Economic Studies* 81(1): 84–111.
- Doeringer, Peter, and Michael Piore. 1971. *Internal Labor Markets and Manpower Analysis*, Lexington, MA: D.C. Heath and Company.
- Fishback, Price, and Joseph A. Cullen. 2013. “Second World War Spending and Local Economic Activity in US Counties, 1939–58.” *Economic History Review* 66(4): 975–992.
- Goldin, Claudia, and Robert Margo. 1992. “The Great Compression: The U.S. Wage Structure at Mid-Century.” *Quarterly Journal of Economics* 107(1): 1–34.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. “Identifying Agglomeration Spillovers: Evidence from Million Dollar Plants.” *Journal of Political Economy* 118(3): 536–598.
- Guiso, Luigi, Luigi Pistaferri, and Fabiano Schivardi. 2005. “Insurance within the Firm.” *Journal of Political Economy* 113(5): 1054–1087.
- Hummels, David, Rasmus Jorgensen, Jakob Munch, and Chong Xiang. 2014. “The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data.” *American Economic Review* 104(6): 1597–1629.
- Kline, Patrick, and Enrico Moretti. 2014. “Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority.” *Quarterly Journal of Economics* 129(1): 275–331.
- Lazear, Edward. 1990. “Job Security Provisions and Employment.” *Quarterly Journal of Economics* 105(3): 699–726.
- . 2009. “Firm-Specific Human Capital: A Skill-Weights Approach.” *Journal of Political Economy* 117(5): 914–940.
- Leduc, Sylvain, and Daniel Wilson. 2017. “Are State Governments Roadblocks to Federal Stimulus? Evidence on the Flypaper Effect of Highway Grants in the 2009 Recovery Act.” *American Economic Journal: Economic Policy* 9(2): 253–292.
- Mayer, Thierry, Marc Melitz, and Gianmarco Ottaviano. 2016. “Product Mix and Firm Productivity Responses to Trade Competition.” NBER Working Paper No. 22433. Cambridge, MA: National Bureau of Economic Research.
- Moretti, Enrico. 2010. “Local Multipliers.” *American Economic Review: Papers & Proceedings* 100(2): 1–7.
- Mortensen, Dale, and Christopher Pissarides. 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies* 61(3): 397–415.
- Piketty, Thomas. 2014. *Capital in the Twenty-First Century*. Cambridge, MA: Harvard University Press.
- Serrato, Juan Carlos Suarez, and Philippe Wingender. 2014. “Estimating Local Fiscal Multipliers.” NBER Working Paper No. 22425. Cambridge, MA: National Bureau of Economic Research.
- Song, Jae, Nicholas Bloom, Fatih Guvenen, David Price, and Till von Wachter. 2015. “Firming Up Inequality.” NBER Working Paper No. 21199. Cambridge, MA: National Bureau of Economic Research.
- van Dijk, Jasper Jacob. 2016. “Local Employment Multipliers in US Cities.” *Journal of Economic Geography* 17(2): 465–487.