Return of the Solow Paradox? IT, Productivity, and Employment in U.S. Manufacturing*

Daron Acemoglu
MIT, CIFAR and NBER

David Autor
MIT and NBER

David Dorn
CEMFI and IZA

Gordon H. Hanson
UCSD and NBER

Brendan Price
MIT

December 2013

Still Preliminary and in Progress

Abstract

An increasingly influential "technological-discontinuity" paradigm suggests that IT-induced technological changes are rapidly raising productivity while making workers redundant. This paper explores the evidence for this view in the U.S. manufacturing sector. In contrast to this view and our expectations, we find little differential productivity growth in IT-intensive manufacturing industries. In fact, gross output does not even appear to grow faster in such industries than in the rest of manufacturing. Though there is some relative decline in employment in these industries in the 1990s, this is preceded by more rapid growth in the 1980s and is followed by more rapid growth in the 2000s. Overall, there is little support for this popular technological-discontinuity view within U.S. manufacturing.

Keywords: employment, IT capital, productivity, Solow Paradox.

JEL Classifications: O3, J2, L6.

*Prepared for the *AEA Papers and Proceedings 2014*. Acemoglu and Autor acknowledge financial support from the Sloan Foundation (Grant 2011-10-12). Autor and Hanson acknowledge funding from the National Science Foundation (Grant SES-1227334). Dorn acknowledges funding from the Spanish Ministry of Science and Innovation (ECO2010-16726 and JCI2011-09709). Price acknowledges financial support from the Hewlett Foundation.

An increasingly popular "technological-discontinuity" perspective, powerfully articulated and defended in Brynjolfsson and McAfee (2011), argues that U.S. workplaces have been, and will continue to be, automated and transformed by information technology (IT) capital. Two implications of this transformation are emphasized. First, all sectors—but particularly IT-intensive sectors—are experiencing major increases in productivity. Thus, Solow's paradox is long since resolved: computers are now everywhere in our productivity statistics.[1] Second, IT-powered machines will increasingly replace workers—especially less skilled production workers—ultimately leading to a substantially smaller role for labor in the workplace of the future.[2]

Several forms of indirect evidence substantiate this argument. The wage premium paid to the most skilled and educated workers, those with post-college degrees, has risen substantially in recent decades, not only relative to non-college workers but also relative to those with a four-year college degree (Acemoglu and Autor, 2011). Berman, Bound and Griliches (1994) and Autor, Katz and Krueger (1998) show that the demand for skills increased more rapidly in the 1980s and 1990s in IT-intensive sectors. The more recent growth of employment of post-college educated workers has also been concentrated in the sectors of the economy undergoing the most rapid computerization (Lindley and Machin, 2013). Furthermore, labor's share of national income has fallen in numerous developed and developing countries over roughly the last three decades, a phenomenon that Karabarbounis and Neiman (forthcoming) attribute to IT-enabled declines in the relative prices of investment goods. And many scholars have pointed to the seeming "decoupling" between robust U.S. productivity growth and sclerotic or negligible growth rates of median U.S. worker compensation (Fleck et al. 2011) as evidence that the "race against the machine" has already been run—and that workers have lost.

In this paper, we provide a simple evaluation of this viewpoint using detailed data from the U.S. manufacturing sector. Our findings do not square with either the (positive) productivity side or the (negative) labor-obsolescence side of the technological-discontinuity story. Though we confirm the earlier findings of Brynjolfsson and Hitt (1997), Oliner and Sichel (2000) and Stiroh (2002) that there was faster productivity growth in the mid-1990s in more IT-intensive sectors of the economy (naturally presumed to be those more likely to benefit from computerization), we find that

[1] Robert Solow's comment on computers appears in his 1987 *New York Times Book Review* article: "...what everyone feels to have been a technological revolution, a drastic change in our productive lives, has been accompanied everywhere, including Japan, by a slowing-down of productivity growth, not by a step up. You can see the computer age everywhere but in the productivity statistics."

[2] Of course, there have also been acerbic critiques of this view, for example by Herbert Simon who as far back as 1966 wrote in the *New York Review of Books*: "...the world's problems in this generation and the next are problems of scarcity, not of intolerable abundance. The bogeyman of automation consumes worrying capacity that should be saved for real problems—like population, poverty, the bomb, and our own neuroses."

this pattern has not persisted. Excluding the computer-producing sector (NAICS code 334, which notably encompasses semiconductor manufacturing as well as computers themselves), we find no indication of a recent (post-2000) or ongoing IT-driven productivity acceleration. In fact, perhaps more surprisingly, we find no evidence of a more rapid increase in shipments or value-added in more IT-intensive industries—a minimum requirement if computers and automation are indeed reducing costs and increasing productivity in the IT-intensive manufacturing sectors (relative to the rest of manufacturing).

Employment in IT-intensive industries offers a similarly counterintuitive pattern. Contrary to the vision of the worker-less workplace, payroll employment and wage bills actually grew relatively faster in IT-intensive industries than in the rest of manufacturing up to the mid-1980s. Though IT-intensive industries have experienced slower employment and wage-bill growth since 1990, cumulative growth of these outcomes for both production and non-production workers has been only modestly slower in IT-intensive industries between the late 1970s and the present.

Our results have several possible explanations. First, our results may be an artifact of our measurement strategy. We discuss this concern and provide evidence that our strategy does not unduly bias our results towards finding an end to the impact of computers on productivity and employment after the late 1990s. Second, our results may be a consequence of our focus on manufacturing. Though this is possible, it is made less likely by earlier evidence from Stiroh (2002) which suggests that the IT-driven productivity productivity growth in the 1990s was not specific to non-manufacturing and may have in fact been more pronounced in manufacturing. Third, it may well be that computers are transforming the U.S. workplace, but this is not related in any way to IT investments and IT intensity. Though this is not entirely implausible, we are not aware of any direct evidence supporting this possibility. Overall, our results suggest (at least to us) that although computerization has demonstrably altered the organization and content of work in recent decades, it has not so far represented a fundamental economic discontinuity relative to the last two centuries of rapid technological progress, as synthesized, for example, by Katz and Margo (2013).

I. Information Technology and Labor Productivity

We employ the NBER-CES Manufacturing Industry Database, sourced primarily from the Annual Survey of Manufacturers, to estimate and plot a set of simple, descriptive regressions that chart the relationship between IT investment and industry-level outcomes for the time period 1977 through 2009. Our regression model takes the form

$$Y_{jt} = \gamma_{\mathbf{j}} + \delta_{\mathbf{t}} + \sum_{t=78}^{09} \beta_t \times IT_j + e_{jt},$$

where Y is an outcome variable, γ is a vector of industry fixed effects, δ is a vector of time dummies, IT is a baseline measure of industry IT-intensity, and e is an error term. This specification normalizes the coefficient on the IT variable to zero in the base year, and hence the series $\{\beta_{78}, \beta_{79}, ..., \beta_{09}\}$ may be read as the level of the coefficient on IT in each subsequent year relative to 1977. Following Berman, Bound and Griliches (1994) and Autor, Katz and Krueger (1998), we measure IT intensity as the ratio of industry computer (IT) expenditures to total capital expenditures.[3]

Figure 1A, which plots the over-time relationship between IT-intensity and the log of real shipments per worker (our preferred productivity measure),[4] shows a dramatic differential rise in output per worker in IT-intensive industries throughout the entire 1977 through 2009 period.[5] But crucially, this pattern is almost entirely driven by the computer-producing sector (NAICS 334).[6] Across the entire manufacturing sector, industries that had a one standard deviation higher rate of IT investment over the first ten years of the sample saw differential productivity gains averaging a remarkable 15 to 20 log points per decade between 1977 and 2009. Excluding the computer-producing sectors, however, results in a far less clear picture. There is some differential growth of IT-intensive industries in the late 1990s, but this is very small (on the order of a few percentage point at its peak) and subsides after 2001. By 2009, there is no net relative productivity gain in IT-intensive industries over the full sample period.[7]

This productivity and output growth pattern is unexpected in light of the earlier resolution of the Solow Paradox. One possible explanation is that our focus on manufacturing is misplaced—perhaps

[3] Specifically, we compute this ratio in 1977, 1982, and 1987, take the average across these three years, and standardize the result so that the final measure has zero mean and unit standard deviation across employment-weighted industries.

[4] A key advantage of this productivity construct, relative to value-added per worker, is that it is unaffected by the choice of deflators for intermediate inputs; so, if the productivity of a dollar of IT investment rises over time due to improvements in IT quality, this will raise shipments in IT-using sectors. By contrast, the effect of rising IT quality on value-added and TFP in IT-using sectors is ambiguous. If these quality improvements were fully captured by deflators (as should ideally occur), the productivity gains would be fully attributed to the IT-producing rather than the IT-using sectors. For completeness, we also show results with value added below.

[5] The dotted lines depict 95 percent confidence intervals using robust standard errors that are clustered at the industry level.

[6] Our focus on NAICS 334 follows Houseman et al. (2013). Because our sample period straddles the transition from SIC to NAICS codes in 1997, we crosswalk data from all years into a consistent set of 387 4-digit manufacturing industries closely based on the SIC classification. The 28 industry codes comprising our computer sector constitute virtually a one-to-one correspondence with NAICS 334.

[7] Houseman et al. (2013), along with other authors, have underscored that the relatively robust growth of productivity and value-added in U.S. manufacturing over the last two decades is substantially driven by IT-producing sectors. What Figure 1 contributes to this discussion is the finding that outside of the IT-producing sectors, there is little relationship between IT investments and productivity growth.

the productivity gains from IT investments are taking place elsewhere. Two arguments weigh against this. One is that productivity growth in U.S. manufacturing has generally exceeded that outside of manufacturing for many decades, and this productivity growth differential rose sharply during the 1990s (Fleck et al., 2011). A second is that Stiroh's (2002) analysis of the post-1995 IT-driven rise in productivity relative to the preceding decade is far from being confined to non-manufacturing.[8]

A second possibility is that our measure of IT investment, based on data from 1977, 1982 and 1987, misses the mark. While we view it as advantageous that this measure potentially captures a durable (or "permanent") aspect of industry IT intensity, it could simply be that past measures of IT investment do not reflect the current locus of the IT frontier. One way to explore this conjecture is to ask whether more recent measures of IT investment perform better than more distant measures in predicting industry productivity growth. Panel B of Figure 1 performs this test by plotting the over-time relationship between labor productivity in non-computer-producing industries and our three different vintages of IT investment measures (one each from 1977, 1982 and 1987) as well as our preferred measure, which is simply the arithmetic mean of the three. (In this panel, we suppress the standard error bands to increase legibility). This simple analysis does not lend support to the hypothesis that the IT measures are simply "out of date." Indeed, the strongest predictor of industry relative productivity growth during the 1990s and 2000s is the 1982 investment measure, followed by the 1977 investment measure. The most recent measure (from 1987) actually performs worse than the other two.[9] This evidence is not definitive, of course, and we plan to incorporate more recent (post-1987) ASM computer investment measures to further explore these relationships. With the available evidence, our tentative conclusion is that the IT-driven productivity acceleration commencing in the mid-1990s was small and short-lived, at least in manufacturing.

II. What Drives Rising Y/L—The Numerator or the Denominator?

Since our measure of labor productivity is equal to the log ratio of gross output to payroll employment, the positive post-1995 relationship we detect in Figure 1A between industry IT investment and output per worker implies that industry output is rising proportionately faster than employment in IT-intensive industries, but does not reveal whether either output is rising faster or employment is falling more rapidly relative to non-IT-intensive sectors. The same applies to the negative post-2001

[8] In fact, columns 4 and 5 in the first panel of Table 2 of Stiroh (2002) show somewhat slower differential productivity growth of IT-intensive industries relative to 1987-1995 when durable goods manufacturing is excluded from the sample, though the pattern is reversed when the comparison is to 1977-1995 in the second panel.

[9] All IT investment measures are standardized with mean zero and variance one. The unstandardized values are, in percentage points: 1977, mean (SD) 2.32 (3.7); 1982, mean (SD) 3.47 (3.77); 1987, mean (SD) 6.84 (6.04); and 1977/1982/1987 average, mean (SD) 4.21 (4.0).

relationship also visible in Figure 1A.

We explore these two outcomes (output and employment) separately in the next two figures. Under the assumption that IT-intensive industries are seeing improvements in technology and automation and reductions in production costs, we would expect them to experience a relative expansion, particularly after 1995. The implications for employment are ambiguous—and this could make the labor productivity measure somewhat more difficult to interpret—because these sectors may be shedding labor as they automate (as claimed by the strong "race against the machine" view) but may also increase their employment as they expand.

The first panel of Figure 2 examines the numerator of this ratio, the logarithm of (real) shipments, while the second considers an alternative measure of industry output, the logarithm of (nominal) value-added.[10] We exclude the computer-producing sector from this analysis (and all subsequent analyses), since our focus is on *induced* productivity gains in IT-using sectors. The relationship between IT-intensity and industry output is almost precisely the opposite of our expectations: both shipments and value-added rise differentially and then plateau in IT-intensive industries through the early 1990s—during the heyday of the Solow Paradox—and then fall relative to other sectors in the period during which the IT productivity payoff is thought to have materialized. Indeed, shipments and value-added in IT-intensive industries fall by 20 to 30 log points relative to other manufacturing industries (per standard deviation unit of IT investment) between 1991 and 2004, and then begin modest rebounds in the mid-2000s—though they both remain below their 1977 starting points at the end of the sample. While it is possible to argue that demand for the output of IT-intensive industries is price inelastic, this would not explain the nominal value-added results.

The differential growth of total, production and non-production employment and wage bill by IT intensity is more nuanced as shown in Figure 3. Employment and wage bill expand more rapidly in IT-intensive industries until the late 1980s, and thereafter start contracting. This contraction reaches its nadir in 2002. Thereafter, both employment and wage bill in IT-intensive industries rebound relative to the rest of manufacturing. At the end of the sample, both are only slightly below their starting value in 1977.[11] This, overall, shows some declines in employment (and labor demand as measured by the wage bill) in IT-intensive industries consistent with the "race against the machine" view, but the extent of this effect is limited and its pattern is variegated.

[10] We are currently calculating a chain-weighted real value added measure.

[11] In a long-distance regression employment and wage bill growth between 1977 and 2009 are not statistically significantly related to IT intensity.

III. Conclusion

This paper documents a pattern of growth among manufacturing industries that stands in contrast to the powerful and intuitively-appealing technological-discontinuity view that IT-induced technological changes are transforming the modern workplace and making workers redundant. We find little differential productivity growth in IT-intensive manufacturing industries. In fact, such industries do not even appear to experience faster growth of gross output than the rest of manufacturing. Though there is some relative decline in employment in these industries in the 1990s, this follows their more rapid growth in the 1980s and is followed by their relative rebound in the 2000s. Overall, there is little support for this popular technological-discontinuity view within U.S. manufacturing.

We view our paper not as a definitive refutation of this influential view, but as a puzzle that needs to be addressed by more systematic empirical work in the future. It may well be that IT-induced technological changes are transforming non-manufacturing, or that they are so widespread as to be taking place rapidly even in non-IT-intensive industries. But at the very least, our evidence suggests that the previously-proposed resolutions of the Solow Paradox need to be critically examined and proponents of the technological-discontinuity view need to provide more evidence of the IT-induced transformation in U.S. economy.

References

Acemoglu, Daron and David H. Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics, Volume 4*, Amsterdam: Elsevier-North Holland, 2011, 1043–1171.

Autor, David H., Lawrence F. Katz and Alan B. Krueger, 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics*, 113 (4), November, 1169–1214.

Berman, Eli, John Bound, and Zvi Griliches. 1994. "Changes in the Demand for Skilled Labor Within U.S. Manufacturing Industries: Evidence from the Annual Survey of Manufacturing." *Quarterly Journal of Economics*, 109(2), 367–397.

Brynjolfsson, Erik and Lorin Hitt. 1996. "Paradox Lost? Firm-Level Evidence on the Returns to Information Systems Spending." *Management Science*, 42(4).

Brynjolfsson, Erik and Andrew McAfee. 2011. *Race Against the Machine.* Lexington, MA: Digital Frontier Press.

Fleck, Susan, John Glaser, and Shawn Sprague. 2011. "The Compensation-Productivity Gap:

A Visual Essay.” *Monthly Labor Review*, January, 57-69.

Houseman, Susan N., Timothy Bartik, and Timothy Sturgeon. 2013. “Measuring Manufacturing: Problems of Interpretation and Biases in the U.S. Statistics.” Mimeo, Upjohn Institute.

Karabarbounis, Loukas and Brent Neiman. Forthcoming. “The Global Decline of the Labor Share.” *Quarterly Journal of Economics.*

Katz, Lawrence F. and Robert A. Margo. 2013. “Technical Change and the Relative Demand for Skilled Labor: The United States in Historical Perspective.” NBER Working Paper No. 18752, February.

Lindley, Joanne and Stephen Machin. 2013. “The Rising Postgraduate Wage Premium.” University College London Working Paper, October.

Oliner, Stephen D. and Daniel E. Sichel. 1994. “Computers and Output Growth Revisited: How Big is the Puzzle?” *Brookings Papers on Economic Activity*, 2, 273-317.

Oliner, Stephen D. and Daniel E. Sichel. 2000. “The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?” *Journal of Economic Perspectives*, 14(4), 3-22.

Simon, Herbert H. 1966. “Automation. In Response to Where Do We Go From Here?” *The New York Review of Books*, March 17.

Solow, Robert. 1987. “We’d better watch out.” *New York Times Book Review*, July 12, p 36.

Stiroh, Kevin J. 2002. “Information Technology and the U.S. Productivity Revival: What Do the Industry Data Say?” *American Economic Review*, 92(5), 1559–1576.

Figure 1A: Effect of IT Intensity on Log Real Shipments/Worker

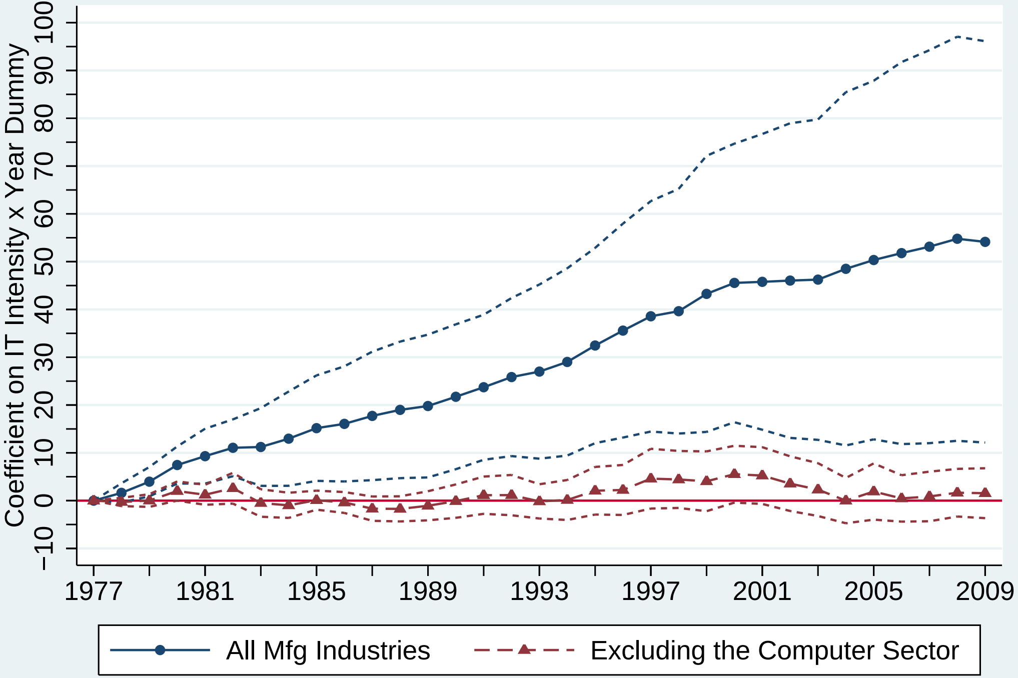

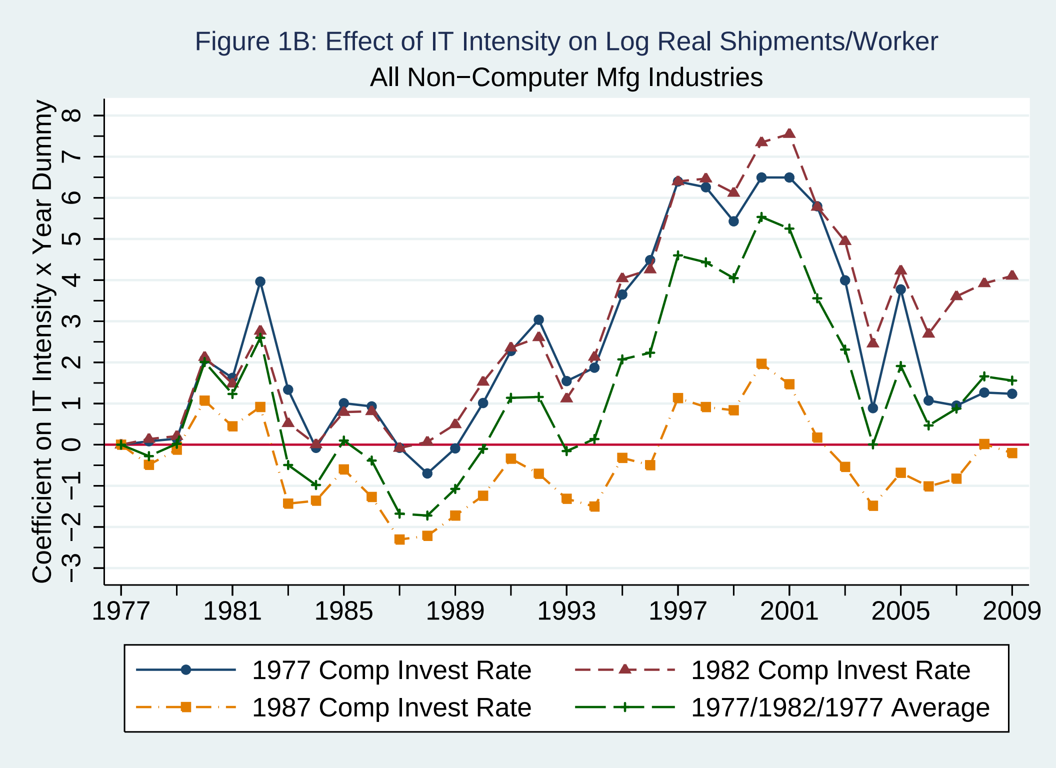

Figure 1B: Effect of IT Intensity on Log Real Shipments/Worker
All Non−Computer Mfg Industries
Coefficient on IT Intensity x Year Dummy
−3
−2
−1
0
1
2
3
4
5
6
7
8
1977
1981
1985
1989
1993
1997
2001
2005
2009
1977 Comp Invest Rate
1982 Comp Invest Rate
1987 Comp Invest Rate
1977/1982/1977 Average

Figure 2A: Effect of IT Intensity on Log Real Shipments

All Non−Computer Mfg Industries

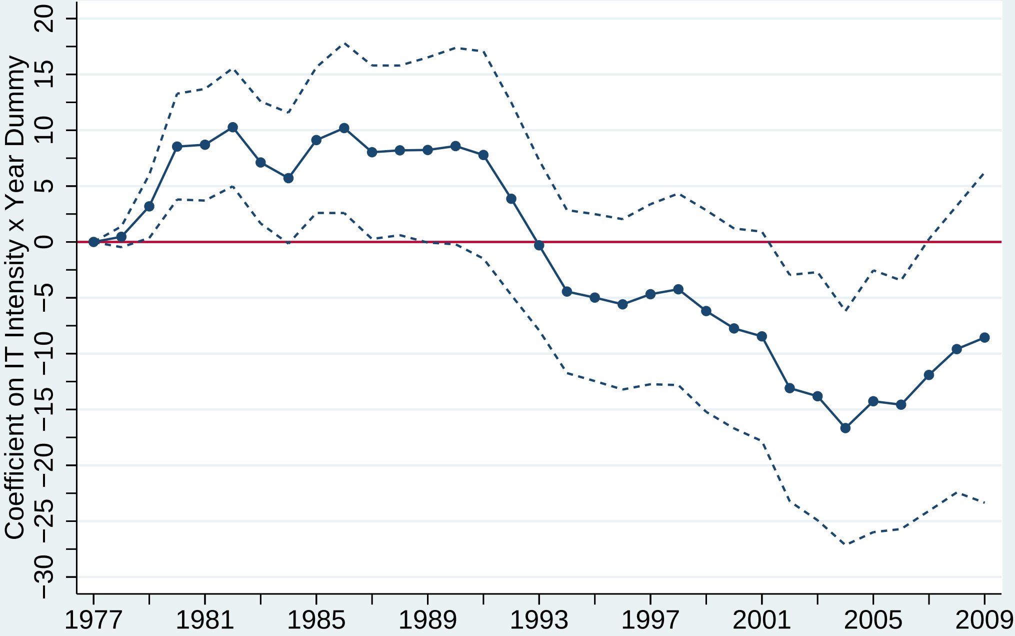

Figure 2B: Effect of IT Intensity on Log Nominal Value Added

All Non−Computer Mfg Industries

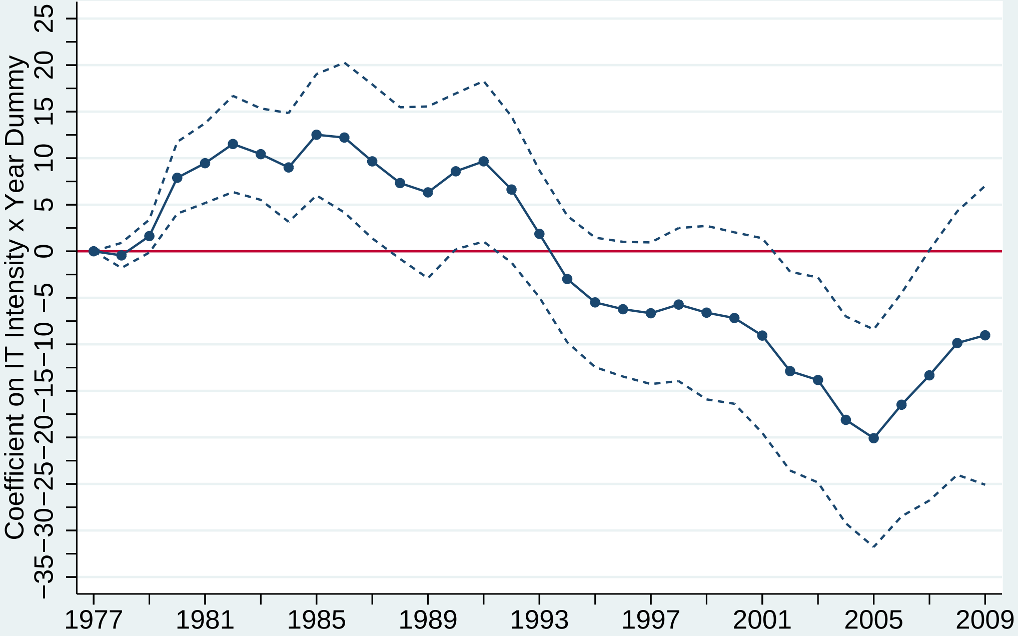

Figure 3A: Effect of IT Intensity on Log Employment and Log Real Wage Bill

All Non−Computer Mfg Industries

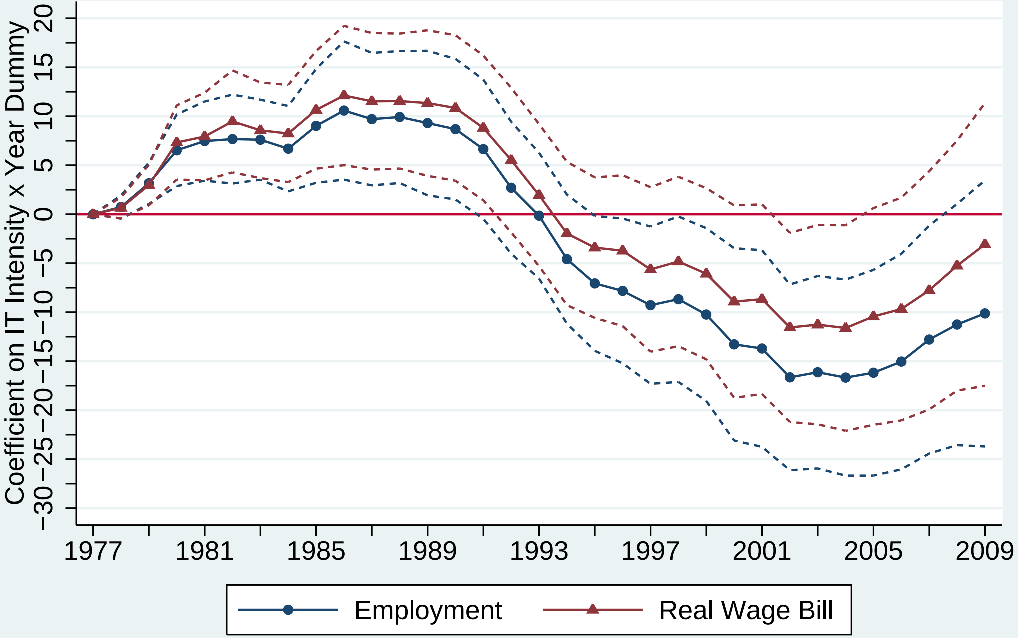

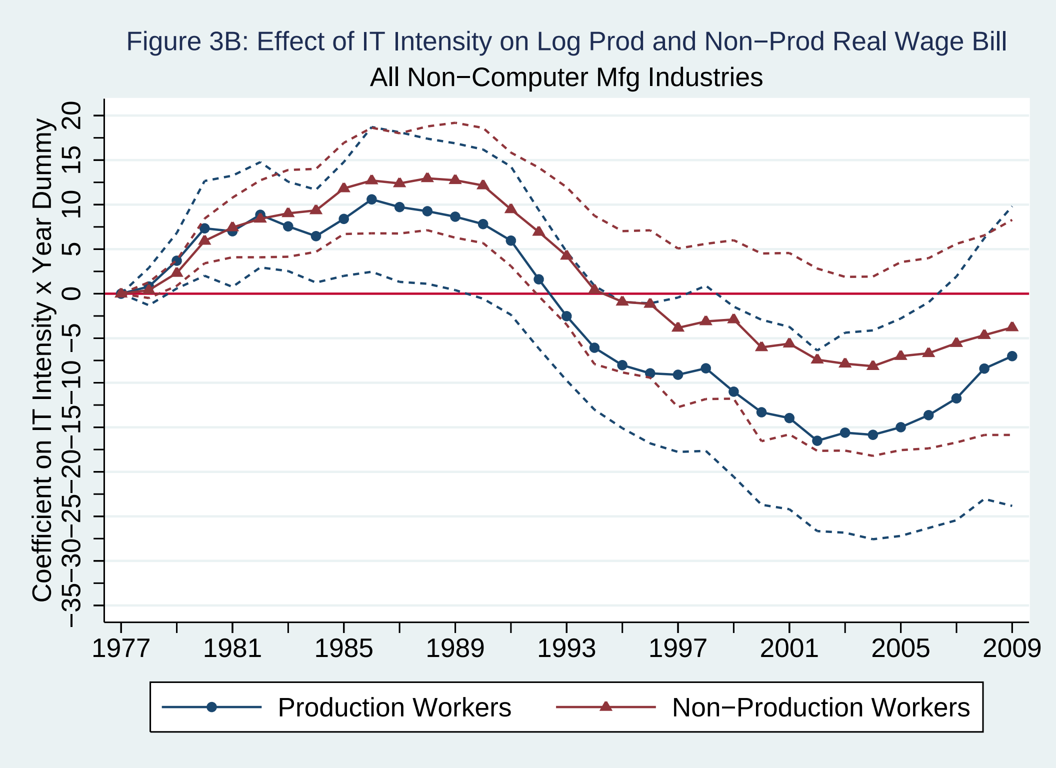
Figure 3B: Effect of IT Intensity on Log Prod and Non−Prod Real Wage Bill
All Non−Computer Mfg Industries
Coefficient on IT Intensity x Year Dummy
20
15
10
5
0
−5
−10
−15
−20
−25
−30
−35
1977
1981
1985
1989
1993
1997
2001
2005
2009
Production Workers
Non−Production Workers