A Short-Run View of What Computers Do: Evidence from a U.K. Tax Incentive

August 16, 2016

Abstract: We study the short-run, causal effect of Information and Communication Technology (ICT) adoption on the employment and wage distribution. We exploit a natural experiment generated by a tax allowance on ICT investments and find that the primary effect of ICT is to complement non-routine, cognitive-intensive work. We also find that the ICT investments led to organizational changes that were associated with increased inequality within the firm and we discuss our findings in the context of theories of ICT adoption and wage inequality. We find that tasks-based models of technological change best fit the patterns that we observe.

JEL: J24, J31, J82, O33
Keywords: ICT, job polarization, wage inequality, routine-biased technical change, organizational change, tax incentive

Paul Gaggl[1,2]

University of North Carolina at Charlotte
Belk College of Business
Department of Economics
9201 University City Blvd
Charlotte, NC 28223-0001
Email: pgaggl@uncc.edu

Greg C. Wright[1,2]

University of California, Merced
Department of Economics
5200 North Lake Blvd.
Merced, CA 95343

Email: gwright4@ucmerced.edu

[2]For extremly helpful feedback we are grateful to David Autor, Steve Billings, Tom Blake, Craig Depken, Rowena Gray, Nir Jaimovich, Jason Lindo, Giovanni Mastrobuoni, Ankur Patel, Andrea Salvatori, Kurt Schnier, Lisa Schulkind, Caleb Stroup, Alex Whalley and seminar participants at Duke University, the 2014 and 2015 NBER Summer Institutes, the 2015 SOLE Meetings, the 2015 South Carolina Applied Micro Day (SCAM'd), the 2015 All-California Labor Economics Conference, the 2016 AEA Meetings, and the 2016 ZEW conference on "Occupations, Skills, and the Labor Market". We are also indebted to three anonymous referees and the two editors.

1. Introduction

Policymakers and researchers have recently been interested in the relationship between firms' adoption of new information and communication technologies (ICT) and the changing demand for different types of labor. A prominent example is the work by Autor, Levy and Murnane (2003), who address the question of "what computers do" by presenting correlations suggesting that ICT complements work that involves the execution of complex, non-routine workplace tasks, while to an equal or greater extent substituting for work that is highly routine. Substantial long-run evidence consistent with the Autor et al. (2003) thesis has since been presented,[2] yet an important question remains: what is the direct impact of ICT investment on the demand for different types of labor within the firm?

A key challenge in testing the Autor et al. (2003) hypothesis is that over long time horizons the relative supply of different labor types varies in response to improvements in ICT. Furthermore, the opposite is also true—innovations in ICT are driven in part by changes in the skill content of the economy, an argument that has most prominently been advanced by Acemoglu (1998, 2002, 2007). More recently, this two-way interaction was the central theme of work by Goldin and Katz (2008), who document the long-run "race" between education and technology.[3] As a result, in order to isolate the direct relationship between ICT adoption and the demand for different types of labor, it is important to hold both the supply of skill and the economy-wide level of technology fixed.

We address this challenge by exploiting a unique natural experiment generated by an unanticipated, narrowly targeted U.K. tax incentive. In doing so, we exploit a situation in which otherwise (nearly) identical firms face different effective post-tax prices for their ICT investment. In contrast, the bulk of the existing

[2] For instance, Autor et al. (2003) present correlations between the use of personal computers (PCs) and the prevalence of non-routine work over the period 1960 to 1998, while Akcomak, Kok and Rojas-Romagosa (2013) do so over the period 1997 to 2006. Michaels, Natraj and van Reenen (2014) take an international perspective and report conditional correlations between ICT and labor market outcomes for 11 countries over 25 years. Doms and Lewis (2006), Beaudry, Doms and Lewis (2010), as well as Autor and Dorn (2013) pursue a more causal interpretation by documenting a positive, long-run relationship between the (likely exogenous) historical concentration of routine tasks across local labor markets and subsequent workplace computer adoption. Autor and Dorn (2013) find that these historically routine intensive regions also show rising wages and employment at the tail ends of the skill distribution relative to middle skilled jobs ("job polarization"). Autor, Dorn and Hanson (2013a,b) also pursue this "tasks-based" approach to the labor market and adopt a similar indirect supply-side identification strategy. Acemoglu and Autor (2011) as well as Draca, Sadun and van Reenen (2006) provide recent reviews of this literature.

[3] Another potentially important long-run general equilibrium effect is suggested by Autor and Dorn (2013). They argue that the observed increase in the demand for low-skill, non-routine-intensive service work (e.g., health care, massage therapy, food service, etc.) throughout the 1990s and 2000s could be a consequence of ICT adoption. Specifically, they argue that the increased demand for these services may come from workers who are complementary to ICT and have therefore seen their incomes, and marginal propensity to consume services, rise.

research that attempts to estimate the causal effect of ICT has typically assumed that the same technology is available to all firms, industries or cities (at the same time and at the same price) but that some of these firms are more likely than others to adopt the technology due to pre-determined variation in workforce composition. Our focus on a shock to the price of ICT investment for some, but not all, otherwise identical firms arguably allows for a more direct estimate of the firm-level impact of ICT investment.

Since exogenous variation in ICT investment is difficult to find, causal evidence on the direct consequences of investing in these technologies is rare.[4] Two analyses that are closely related to ours are Akerman, Gaarder and Mogstad (2015) and De Stefano, Kneller and Timmis (2014). The former exploits the sequential rollout of broadband internet across Norway and finds a significant impact on firm performance as well as the wage distribution within firms. In contrast, the latter exploits a geographic discontinuity in the availability of broadband internet in the U.K. but finds little effect on firm performance. Our paper is similar to this work in that we also exploit a natural experiment in order to generate firm- and worker-level estimates. In addition, the paper is in line with a related literature that proxies ICT adoption with either the number of workers that uses a PC or the number of PCs in the workplace (e.g., see Beaudry et al., 2010, Autor and Dorn, 2013, and much of the literature surveyed in Acemoglu and Autor, 2011).

Our research design exploits a 100 percent tax credit for investments in ICT that was made available exclusively to small U.K. firms between 2000 and 2004.[5] Since the tax incentive was largely unanticipated and, in addition, was the only tax policy targeted exclusively toward small firms at the time, it generated a sharp discontinuity in the incentive to invest in ICT at the eligibility threshold.[6] Thus, any observed

[4]While many of the existing contributions only present conditional correlations, attempts to identify the causal effects of ICT on workers go back at least to Krueger (1993) and DiNardo and Pischke (1997). More recent papers that attempt to identify causal effects include Agha (2014), Ivus and Boland (2015), Bertschek, Cerquera and Klein (2013), Atasoy (2013) or see Cardona, Kretschmer and Strobel (2013) for a review of the recent literature.

[5]Our analysis complements a handful of existing studies on the short-run consequences of ICT adoption. These are based on less comprehensive datasets and have a slightly different focus. For instance, Bartel, Ichniowski and Shaw (2007) use survey responses from 212 US valve-making plants in 2002 to study potential plant-level mechanisms through which computers may enhance productivity. Among other results, they find that the "adoption of new IT-enhanced capital equipment coincides with increases in the skill requirements of machine operators, notably technical and problem-solving skills, and with the adoption of new human resource practices to support these skills." In a related paper, Brynjolfsson and Hitt (2003) study 527 large US firms over 1987-1994 and provide evidence for overall short-run firm-level productivity gains in response to computer investments, though they do not study the effect of computers on labor demand. Doms, Dunne and Troske (1997) focus on the manufacturing sector and present industry-level correlations that suggest that computers increase demand for non-production labor.

[6]All other corporate tax incentives applicable to small firms (with no more than 50 employees) at the time were also available to medium sized firms (with no more than 250 employees). Interestingly, Abramovsky and Griffith (2006) use the same tax incentive as we do to construct an instrumental variable in their study of services offshoring.

discontinuity in firm investment behavior at this threshold can potentially be interpreted as a causal outcome of the policy. With this in mind, we follow a standard regression discontinuity (RD) design throughout, first exploiting data from the universe of U.K. corporate tax returns to show that the introduction of this targeted tax credit differentially altered tax claims for ICT investment at the eligibility threshold. Using separate firm-level data on investments in computer hardware and software we then confirm that the additional tax claims made by small firms indeed reflected actual additional investment in ICT.

Shifting our focus to workers, we find that these additional ICT investments raised average weekly earnings, hours worked, and labor productivity within the treated group of firms. In order to explore the distributional consequences of the ICT investments we adopt the job-type classification from Acemoglu and Autor (2011), finding that the rise in earnings and employment was concentrated among workers engaged in non-routine, cognitive-intensive production tasks.[7] On the other hand, we find that routine, cognitive workers experienced a decline in earnings and employment while there was little evident impact on manual workers. Overall, this asymmetric pattern is consistent with an outward shift in the demand for non-routine, cognitive-intensive production tasks and an inward shift in the demand for routine, cognitive work in response to the investment in ICT. Finally, we find that the dispersion in wages among non-routine, cognitive workers increased in response to the ICT investments, while there is some evidence that wage dispersion fell among routine, cognitive workers. We later highlight this finding in our discussion of potential theoretical explanations of the results.

It is not immediately obvious why we should expect a pronounced short-run labor demand response to ICT investments. One possible explanation is that organizational changes are implemented as part of the ICT-adoption process, and these organizational changes amplify the impact of ICT. For instance, new ICT may allow firms to introduce new business processes, restructure workforce hierarchies, or monitor workers more or less intensively, which may consequently affect the productivity of different categories of workers. While the existing literature has explored the potential complementarity of ICT and organizational change, the firm's organizational status is typically either assumed to be exogenous[8] or else it is endogenized but

[7]Following Acemoglu and Autor (2011), we classify (1) managerial, professional and technical occupations as "non-routine, cognitive"; (2) sales, clerical and administrative support occupations as "routine, cognitive"; (3) production, craft, repair, and operative occupations as "routine manual"; and (4) service occupations as "non-routine manual".

[8]For instance, Bartel et al. (2007), Black and Lynch (2001), and Bresnahan, Brynjolfsson and Hitt (2002).

its relationship to ICT investments is not cleanly identified.[9] We bring our research design to bear on these issues by applying our RD strategy to firm-level data on the implementation of organizational changes. Interestingly, we find evidence that the ICT investments were indeed accompanied by an increase in the adoption of advanced management techniques as well as additional changes in organizational structure.

Taken together with the wage and employment results, we conclude with a discussion of the extent to which the findings align with the recent tasks-based model of technological change presented in Acemoglu and Autor (2011) as well as the model of knowledge hierarchies developed in Garicano and Rossi-Hansberg (2006). There has been some debate over the relative importance of each mechanism in driving recent trends in inequality; for instance, Garicano and Rossi-Hansberg (2015) argue that in contrast to the labor literature, which focuses narrowly on the relative price of skill, their knowledge-based hierarchies model can more readily explain the recent pattern of U.S. wage polarization. With this in mind, we first argue that both models present mechanisms that are consistent with our results. In an attempt to distinguish between the models we present additional findings indicating little impact of ICT on managerial wages, which are a subset of the non-routine, cognitive jobs we examined previously. Since the Garicano and Rossi-Hansberg (2006) model predicts increased wage dispersion among this group of workers, we suggest that this finding points toward a more prominent role for the tasks-based model relative to the knowledge hierarchy model within the particular context of this natural experiment.

Finally, we note that temporary tax incentives, like the one explored here, are a popular vehicle to promote investment. However, there is little convincing evidence on the efficacy of such policies. Our paper therefore contributes to a long debate that goes back at least to Hall and Jorgenson (1967), and is most closely related to two recent contributions that exploit similar tax incentive programs, Cohen and Cummins (2006) and House and Shapiro (2008).[10]

The remainder of this article is organized as follows: we begin with a detailed description of the tax policy and our data sources in Section 2. We introduce our research design and discuss its validity in Section 4. Section 5 covers our results. Section 5.1 presents estimates of the tax incentive's impact on firm

[9] For example, see Acemoglu, Aghion, Lelarge, Van Reenen and Zilibotti (2006), Caroli and Reenen (2001), or Colombo and Delmastro (2004).

[10] For other representative contributions to this debate see Auerbach and Hassett (1991, 1992), Cummins, Hassett and Hubbard (1994, 1996), Goolsbee (1998), or Chirinko, Fazzari and Meyer (1999).

investment decisions, Section 5.2 our estimates of the impact of ICT on labor market outcomes, and Section 5.3 our analysis of the relationship between ICT investment and organizational change. Section 6 discusses the results in the context of current models of ICT and wage inequality. We offer some concluding remarks in Section 7.

2. Policy Experiment & Data Sources

We exploit a unique quasi-experiment due to a 100 percent first year tax allowance (FYA) on ICT investments made available to small firms in the U.K. This policy represented a particularly large investment incentive, as it allowed businesses to write off the full cost of ICT investments against their taxable profits.[11] The following types of investments were eligible for the tax allowance:[12]

- Computer equipment comprising computers (ranging from small palmtop organizers to large systems), computer peripherals such as keyboards, printers etc; cabling and other equipment to link computers to each other, or to data networks such as the internet; and dedicated electrical systems for computers.
- High-tech communications technologies comprising WAP (wireless application protocol) phones, 3rd generation (3G) mobile phones and equipment with similar applications and functionality; and set-top boxes that are connected to televisions and are capable of receiving and transmitting information from and to data networks such as the internet.
- Software for use with computers or high-tech communications technologies. This covers all computer software, including new software for use on computers bought before April 1, 2000 and the costs of creating web sites.

To identify the impact of the policy, we exploit both the timing of its introduction as well as its targeted nature. The tax incentive was introduced on April 1, 2000, was initially scheduled to expire on March

[11] As a comparison, over the period 1998 to 2008 there was one other plant and machinery allowance available to small and medium sized firms (up to 250 employees). The rate varied over this period but was set at either 40 or 50 percent. As a further reference, in a recent paper Bøler, Moxnes and Ulltveit-Moe (2012) find a substantial response by Norwegian firms to a 20 percent tax credit on R&D investments. Overall we believe that the tax credit we examine here represented a significant incentive.

[12] See http://www.hmrc.gov.U.K./manuals/camanual/CA23130.htm for official documentation.

31, 2003, but was then extended until March 31, 2004. The tax incentive was further restricted to "small businesses" which Her Majesty's Revenue and Customs (HMRC) defined as those satisfying at least two of the following criteria: annual revenue of no more than £2.8 million, total assets of no more than £1.4 million, and no more than 50 employees.[13] We note that employment is overwhelmingly the key criterion—i.e., there are very few firms with fewer than 51 employees but more than £2.8 million in revenue or £1.4 million in assets. For the purposes of our research design, it is also important to note that during this period there were no other tax allowances targeting small firms exclusively.[14]

We provide direct evidence on the magnitude of investments made by firms in response to the incentive by exploiting confidential tax return data from HMRC as well as data on investments in hardware and software collected by the U.K. Office of National Statistics in the Quarterly Capital Expenditure Survey (QCES).[15] With respect to the first data source, we exploit the universe of plant and machinery tax allowances claimed by U.K. firms in their corporate tax returns, both during (2000 - 2004) and after the period in which the policy of interest was in place (2005 - 2007).[16] While these tax returns do not tell us whether the firm claimed the specific first year allowance we are interested in (the ICT allowance), they do report the sum of all first year plant and machinery allowances claimed by each U.K. firm in a year. We therefore exploit the discontinuity in firm eligibility for the tax incentive to gauge the magnitude of ICT tax allowance claims made in response to it, as we describe further below. Since eligibility is based on firm employment, we merge the tax return data with firm information from Bureau van Dijk's FAME database, which reports (among many other things) the number of employees in each firm. The FAME data encompass all U.K. firms registered with Companies House, which is nearly all U.K. firms. Panel A of Table 1 presents descriptive statistics for the merged sample.

Since the data on tax allowance claims alone cannot provide conclusive evidence of differential investment across the eligibility threshold, we further explore the direct effect of the incentive on hardware and software investment via the QCES, which collects quarterly data on these expenditures from a stratified

[13]For the remainder of this article we use the term "small firm" as a synonym for the official HMRC definition.

[14]There were five other plant and machinery tax allowances available at different points during this period, however each targeted small and medium sized firms (those with fewer than or equal to 250 employees) rather than small firms only. These allowances targeted investments in energy saving capital (from 2001), cars with low carbon dioxide emissions (from 2002), capital for gas refueling stations (from 2002), capital for ring fence trade (from 2002), and environmentally friendly capital (from 2003).

[15]Note that unfortunately it is not possible to link any of the datasets used in the paper, so we rely on independent data sources.

[16]This information is not currently available to researchers for the pre-2000 period.

Table 1: Data Sources & Summary Statistics

	All Firms 2001-2004			≤50 Employees 2001-2004		
	N	Mean	Std. Dev.	N	Mean	Std. Dev.
A. U.K. Corporate Tax Return Data & Bureau van Dijk's FAME						
P&M Tax Claims ($000')	206481	214.38	4140.62	132516	104.02	1849.05
B. Quarterly Capital Expenditure Survey (QCES)						
Labor Productivity	1005109	114.08	878.96	858757	109.97	751.22
Firm Age	199950	9.91	8.53	184986	9.91	8.52
Investment						
Software	1005718	750.73	4306	859350	765.95	4289.47
Hardware	1005718	54.35	1099.79	859350	46.92	981.6
Other Capital	1005718	7148.88	29713.89	859350	6761.7	27304.14
Land & Buildings	1005718	2703.43	19826.05	859350	2091.04	10765.68
Vehicles	1005718	859.19	9135.09	859350	820.95	8455.46
C. Annual Survey of Household Earnings (ASHE)						
Weekly Wages	2712407	356.49	330.34	534547	323.39	299.93
Weekly Hours	2696017	33.75	12.05	532934	34.49	12.8
Non-Routine Cognitive						
Weekly Wages	1014169	557.83	439.49	176392	516.03	419.86
Weekly Hours	1007842	34.85	9.47	176047	36.19	10.02
Wage Disp. (Log 90th - 10th)	1014169	1.3	0.27	176392	1.33	0.29
Routine Cognitive						
Weekly Wages	701433	235.33	147.09	128645	216.19	142.62
Weekly Hours	698220	30.49	11.03	128193	29.18	11.66
Wage Disp. (Log 90th - 10th)	701433	1.23	0.28	128645	1.25	0.29
Non-Routine Manual						
Weekly Wages	411721	260.12	148.92	113906	257.55	152.35
Weekly Hours	408669	34.79	13.06	113571	36.48	12.79
Wage Disp. (Log 90th - 10th)	411721	1.19	0.21	113906	1.2	0.23
Routine Manual						
Weekly Wages	209433	280.52	124.16	47712	259.46	121.98
Weekly Hours	208893	41.82	11.88	47580	41.95	13.01
Wage Disp. (Log 90th - 10th)	209433	1.26	0.27	47712	1.27	0.27
D. U.K. Innovation Survey 2002-2004 (CIS)						
Change in Corporate Strategy	2342	0.17	0.37	1591	0.14	0.35
Adopt Adv. Management Practices	2342	0.16	0.36	1591	0.13	0.33
Change Organizational Structure	2342	0.18	0.39	1591	0.15	0.35
Change Marketing Strategy	2342	0.2	0.4	1591	0.17	0.38

Notes: The data sources are described in the text. Panel A reports raw sample averages. Panel B reports summary statistics of a dataset that was averaged to weighted firm-size-bin means. Panel C reports pooled weighted averages of the indicated survey years. Panel D averages yes/no (coded 1/0) questions concerning the years 2002-2004.

sample of 26,000 to 32,000 firms, and has done so since 2001.[17] In other words, it may be the case that the incentive had no impact on ICT investments, but simply allowed firms below the threshold to claim

[17] Note that the tax incentive was available from 2000 to 2004, so that the QCES data will underestimate the total effect over the period to the extent that there was significant investment made in 2000 in response to the incentive.

investments they would have made anyway. Panel B of Table 1 presents descriptive statistics.

To implement our research design with respect to workers we exploit the U.K. Annual Survey of Household Earnings (ASHE), an annual representative one percent sample of workers drawn from National Insurance records for working individuals. The dataset provides detailed information on the earnings and hours worked of U.K. workers along with basic employment variables such as the detailed industry and occupation of the worker.[18] Importantly, the survey also includes the number of employees associated with each worker's firm, information that we exploit in order to link each worker's employer to the eligibility criteria of the tax incentive. We focus on those workers who work for private companies, since these workers are effectively "treated" by the tax incentive, though we also exploit the public employee data in a placebo test, which we discuss below. It is also worth noting that because the ASHE earnings data is provided by employers, rather than employees, the reported wage and employment values are thought to be more comprehensive and accurate relative to other surveys. The data are available for the period 1997 - 2007 and summary statistics for our worker sample covering the period 2001 - 2004 are provided in panel C of Table 1.

Finally, we exploit data on organizational change from the U.K. Innovation Survey (CIS), the main source of business innovation information for the U.K. The survey asks a stratified random sample of approximately 16,000 firms a series of questions with respect to their innovation activities over a prior three-year period, and in our case we exploit responses for a subset of our treatment period, 2002-2004, as well as the next available post-treatment period survey responses, which covered 2006-2008.[19] Summary statistics for the first wave are provided in panel D of Table 1.

3. Conceptual Framework

A primary focus of the paper is on the extent to which a fall in the effective price of ICT capital leads to changes in wages and employment levels for different types of workers within UK firms. In part, this will depend on whether the subset of firms that responds to the ICT incentive employs a small or large fraction of the potentially impacted workers in the relevant labor market: when the affected firms employ a negligible

[18]Earnings values are deflated using the U.K. CPI.

[19]There was an intermediate survey that covered 2004-2006, which clearly overlaps with the final year of our treatment period and so is not useful as a counterfactual.

share of local employment of these workers then the firms may be price-takers in the labor market, and wages will remain unchanged within the affected firms. Instead, any change in the demand for labor within the firm will manifest via shifts in employment levels. In contrast, when the firms are price-*setters*, the magnitude of the wage effects will depend on multiple factors, including the extent to which local labor markets are segmented, as we discuss further below.

In our case, the set of firms that were eligible for the ICT incentive (those with 50 or fewer employees) comprised 20 percent of total UK employment during the period, though only a subset of those firms actually responded to the incentive (we do not know the number of firms that actually claimed the incentive). This would seem to leave some ambiguity as to whether these firms are indeed price-takers in UK labor markets. However, as we will see, the empirics indicate the existence of wage effects, and as a result our discussion in this section focuses on the case in which firms face an inelastic labor supply.

Before doing so, it is important to note that even when firms are price-takers in the labor market we still may observe changes in workers' wages due to the fact that firms may respond to the ICT incentive by altering the composition of their workforce in ways that are unobservable to the econometrician. For instance, the dataset that we exploit classifies workers by three-digit occupation, and yet firms may respond to the incentive by upgrading or downgrading their workforce within these occupational categories, perhaps by hiring workers with specialized knowledge related to ICT or by reducing the hours worked of employees without this knowledge. As a result, it may seem that wages have adjusted when, in fact, it is the composition of the workforce within the available data categories that has changed.

3.1. When Small U.K. Firms Face an Inelastic Labor Supply

To the extent that there is a wage response, it may arise via different channels. Furthermore, the individual firm responses may generate general equilibrium effects beyond the boundaries of affected firms, and in this subsection we take this into account in formally outlining the economic forces that may be relevant in the context of our research design.

We assume there are J distinct labor markets, each containing $i = 1, ..., I$ firms that produce output Y^i. Production requires $l = 1, ..., L$ types of labor and a single capital type that we take to be (an aggregate of) ICT capital. Denoting X as the $(L + 1) \times 1$ vector of labor and capital supplies in a labor market, factor market clearing in that market is given by

$$X = \sum_{i=1}^{I} c^i(W, r, A^i) Y^i, \tag{1}$$

where $c^i(W, r, A^i)$ is the $(L+1) \times 1$ vector of unit labor requirements plus the unit capital requirement associated with firm i, W is a vector of wages associated with each labor type, r is the effective price of ICT capital, and A^i is a vector of firm production efficiencies. We can totally differentiate (1) and rearrange terms to get

$$\sum_{i=1}^{I} Y^i c_r^i dr = -\sum_{i=1}^{I} Y^i c_W^i dW + dX - \sum_{i=1}^{I} c^i dY^i - \sum_{i=1}^{I} Y^i c_{A^i}^i dA^i, \tag{2}$$

where c_r^i is a vector of derivatives of the unit factor requirements with respect to the effective price of ICT capital, c_W^i is a matrix of cross-price effects associated with each labor type, and $c_{A^i}^i$ is a vector of derivatives of the unit factor requirements with respect to firm efficiency.

The left hand side is the change in the demand for each factor in response to a change in the effective price of ICT capital, while the right hand side reflects firm and labor market responses to the change in factor demand. Specifically, the first term on the right hand side is a matrix of cross-wage effects whose relative magnitude is proportional to the size of the firm; the second term is a vector of changes in the local supply of each factor; the third term is a vector of changes in firm output whose relative magnitude is proportional to the importance of each factor in production; and the fourth term is a vector of within-firm technological changes (for instance, organizational changes) whose relative magnitude is proportional to the size of the firm as well as the importance of each factor in production.

As an example, a fall in the effective price of ICT capital may raise the demand for cognitive-intensive (ICT-using) labor types. This effect must then be offset by some combination of the right hand side (RHS) channels to bring the labor market back to equilibrium. For instance, the wage of cognitive-intensive workers may rise (first term on the RHS), an effect that may then be partially or completely offset by an inflow of cognitive-intensive workers from surrounding regions (second term on the RHS). Furthermore, when there is firm heterogeneity in the importance of each factor in production, cognitive-intensive workers will reallocate toward the most cognitive-labor-intensive firms, raising output (third term on the RHS) in those firms at the expense of other firms—a so-called “scale effect”. This may then be further amplified by complementary investments in the production technology that are relatively greater for the most cognitive-labor-intensive firms (fourth term on the RHS), which will further raise the demand for these workers leading to a (possibly) larger scale effect and a (possibly) higher wage, which may again be partially or completely offset by a further inflow of workers from surrounding regions.

Any wage impact of the ICT investment is therefore moderated by inflows of workers into the labor market. Importantly, this is governed by the degree to which local labor markets are segmented or, alternatively, nationally integrated. When local labor markets are perfectly integrated then national factor markets must also clear, and in this case the return to each factor will be equal across labor markets in equilibrium. The key implication is simply that when there is a change in the effective price of ICT capital, and labor markets are fully integrated, the first term on the right hand side of (2) cannot be the sole force behind the equilibration of factor demand: wages may rise or fall, but this will be mitigated by worker movement such that national factor markets clear.

In light of this analysis, we highlight three results. First, even with fully integrated labor markets a fall in the effective price of ICT capital can alter the distribution of wages in the economy, an effect that may be amplified by firm production and technological responses but will be mitigated by the movement of factors. Conversely, to the extent that labor markets are *segmented*, wage effects may be relatively strong. Second, a fall in the effective price of ICT capital may induce within-firm organizational changes, which we take to be a subset of the technologies represented by A^j. Third, we note that when the categorization of factors is imprecise, changes in the composition of detailed factor types within factor categories may be confused for uniform changes in factor demand within the category. We will return to this framework when interpreting our results.

4. Research Design

Our primary objective is to explore the within-firm impact of policy-induced ICT investments on workers, and we undertake this analysis in Section 5.2. As a first step, we estimate a “first stage” effect in order to quantify the investment response by firms to the tax incentive. In a “second stage” we then estimate the impact of the tax incentive on worker outcomes. Finally, we use the same strategy to investigate the complementary role of organizational change as a co-determinant of the observed effects.

In each case we employ a standard regression discontinuity (RD) research design that exploits firms' differential incentive to invest in ICT in the neighborhood of the 50 employee threshold (Lee and Lemieux, 2010; Imbens and Lemieux, 2008). In the context of the potential-outcomes framework, we denote Y_{it} as the outcome variable of interest—e.g., investment in ICT equipment; $Y_{it}(1)$ as the observed value when firm i is "treated" with eligibility for the ICT tax allowance in period t; and $Y_{it}(0)$ as the observed value if the firm is not treated (i.e., has more than 50 employees). Further denote *Emp*$_{it}$ as the number of employees in firm i at time t. Under the assumption that the treatment is randomly assigned within a small neighborhood around the eligibility threshold, the causal impact of the treatment on the outcome variable for firms within this neighborhood is then given by

$$\begin{aligned}
\tau_{RD} &= E[Y_{it}(1) - Y_{it}(0)|Emp_{it} = 50] \\
&= \lim_{\epsilon \uparrow 0} E[Y_{it}|Emp_{it} = 50 + \epsilon] - \lim_{\epsilon \downarrow 0} E[Y_{it}|Emp_{it} = 50 + \epsilon] \\
&= \lim_{\epsilon \uparrow 0} \mu_-(\epsilon) - \lim_{\epsilon \downarrow 0} \mu_+(\epsilon) \qquad (3)
\end{aligned}$$

We then obtain estimates of τ_{RD} by approximating the regression function to the right, $\mu_+(\epsilon)$, and to the left, $\mu_-(\epsilon)$, of the cutoff with kernel-based local polynomials.[20] Throughout the paper we follow Calonico, Cattaneo and Titiunik (2015) as well as Calonico, Cattaneo and Titiunik (2014) to construct graphical illustrations and statistical estimates of τ_{RD}. For each set of results we report estimates derived from local linear approximations and our inference is based on the bias-correction procedures described by Calonico et al. (2014). In the interest of space, we report the results from local linear specifications throughout the main text. However, we have confirmed that these results are robust to a variety of alternative specifications and these sensitivity analyses are reported in a companion Online Appendix.[21]

[20]Note that, strictly speaking, the limiting arguments in the above RD design are not valid with a discrete running variable like *Emp*$_{it}$, but in practice any observed measurement is discrete. Lee and Card (2008) as well as Lee and Lemieux (2010) discuss the degree to which a discrete running variable may introduce bias.

[21]Moreover, in an earlier version of the paper we have also explored a more traditional difference-in-differences and event study approach (Gaggl and Wright, 2014).

Figure 1: Firm Size Distribution

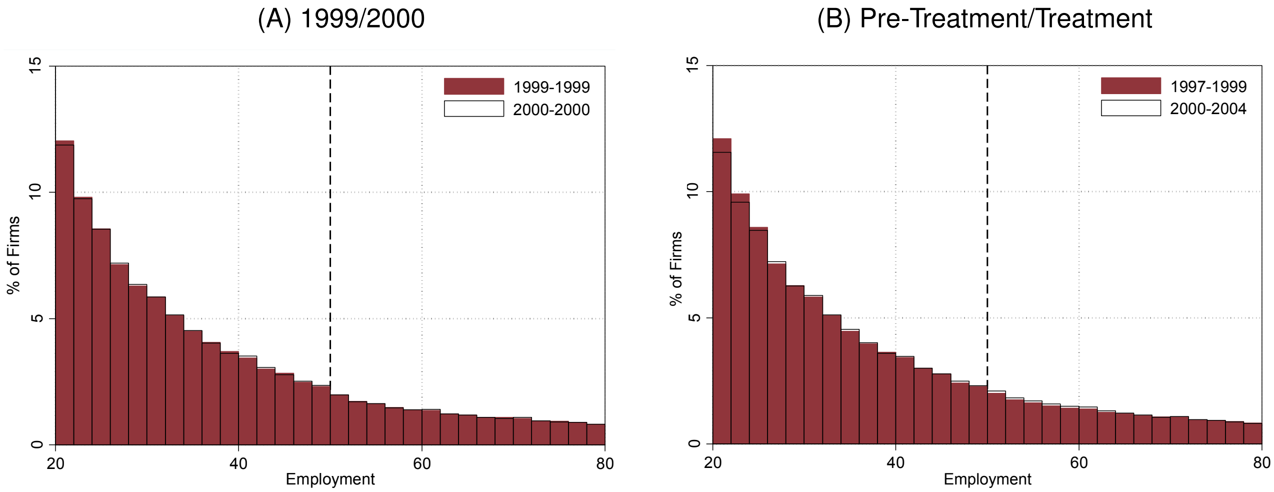

Notes: The figures plot the firm size density before and during the treatment period. Panel A illustrates the densities for 1999 and 2000. Panel B pools the period 1997-1999 and 2000-2004. The data are drawn from Bureau Van Dijk's FAME database.

4.1. Validity of the Research Design

Before turning to the formal analysis outlined above, we first address the validity of our research design and discuss a number of strategies we use to test our identifying assumptions. Since HMRC's definition of a small business was not introduced for this particular policy and since there were no other contemporaneous policies specific to this group of firms, we argue that eligibility for businesses close to the size threshold was effectively randomly assigned. However, to the extent that firms were able to anticipate the tax savings we might expect that firms with initially *more than* 50 employees may have reduced their employment in order to qualify for the incentive. On the one hand, the tax incentive was quite generous, so that the potential tax savings for firms that use ICT intensively may have been substantial. On the other hand, firing workers is costly, and U.K. labor laws are relatively strict. Perhaps a more realistic scenario is that firms who otherwise would have hired additional workers, bringing them over the 50 worker threshold, delayed hiring. These firms may not be otherwise identical to firms with just over 50 employees and the regression discontinuity design would therefore be inappropriate.

Firms who manipulate their size along the lines hypothesized above would likely be concentrated just below the eligibility cutoff in the treatment period and so, as a first pass, we examine the density of firms around the cutoff. Panels A and B in Figure 1 plot the firm size distribution for U.K. firms, drawn from

FAME, for the period 1997-1999 (pre-treatment) and 2000-2004 (treatment), and alternatively for the years 1999 (just before the FYA became available) and 2000 (the first year of the FYA). These distributions suggests that there was no bunching of firms at the 50 employee threshold. Moreover, there is also no visible change in the density before and after the introduction of the ICT tax allowance.

To further assess the role of size manipulation we estimate τ_{RD} using a sample in which we drop observations just below the size threshold for several key specifications—specifically, we drop firms with 48-50 employees. We find that the omission of these firms does not substantively change the results, suggesting that the research design does not simply identify firm selection into eligibility for the tax incentive. These results are reported in the Online Appendix.

Importantly, our research design also relies on the absence of discontinuities at the eligibility threshold for firm outcomes that are exogenous to the tax policy, and in particular outcomes that are pre-determined. In the case examined here it is difficult to identify outcomes that satisfy this criterion for at least two reasons. First, almost all observed firm characteristics are potentially endogenous to the tax incentive. Second, firm entry and exit induced by the tax incentive may affect the estimated cell means, even if the chosen outcome is pre-determined at the level of the firm. For instance, firm age is certainly pre-determined, however its cell mean may be endogenous to the incentive to the extent that the incentive fostered churning in the market. On the other hand, the firm response to the incentive, while of an important magnitude, was not drastic, and particularly over this short time frame is unlikely to have altered the composition of firms to a large extent.

With these caveats in mind, Figure 2 presents regression discontinuity plots of the share of firms in a particular industry by firm size. Figure 3 then plots average firm age and average investment in non-ICT capital by firm size. As expected, there is no evidence of a significant discontinuity in each case.[22] This finding suggests that any discontinuity in ICT investment at the 50 employee threshold is not a general feature of firms around that threshold but, rather, is likely a consequence of the targeted investment incentive.

Of particular interest is the finding of no effect on non-ICT capital investments shown in panel C of Figure 3, which highlights the effectiveness of the tax incentive in targeting specific capital types. In part this is important because it provides evidence against an alternative story in which the ICT tax incentive simply relaxed the liquidity constraint for some subset of small firms, allowing them to make a wide range

[22]This observation is supported by point estimates for this exercise reported in the Online Appendix.

Figure 2: Industry Composition: 2001-2004

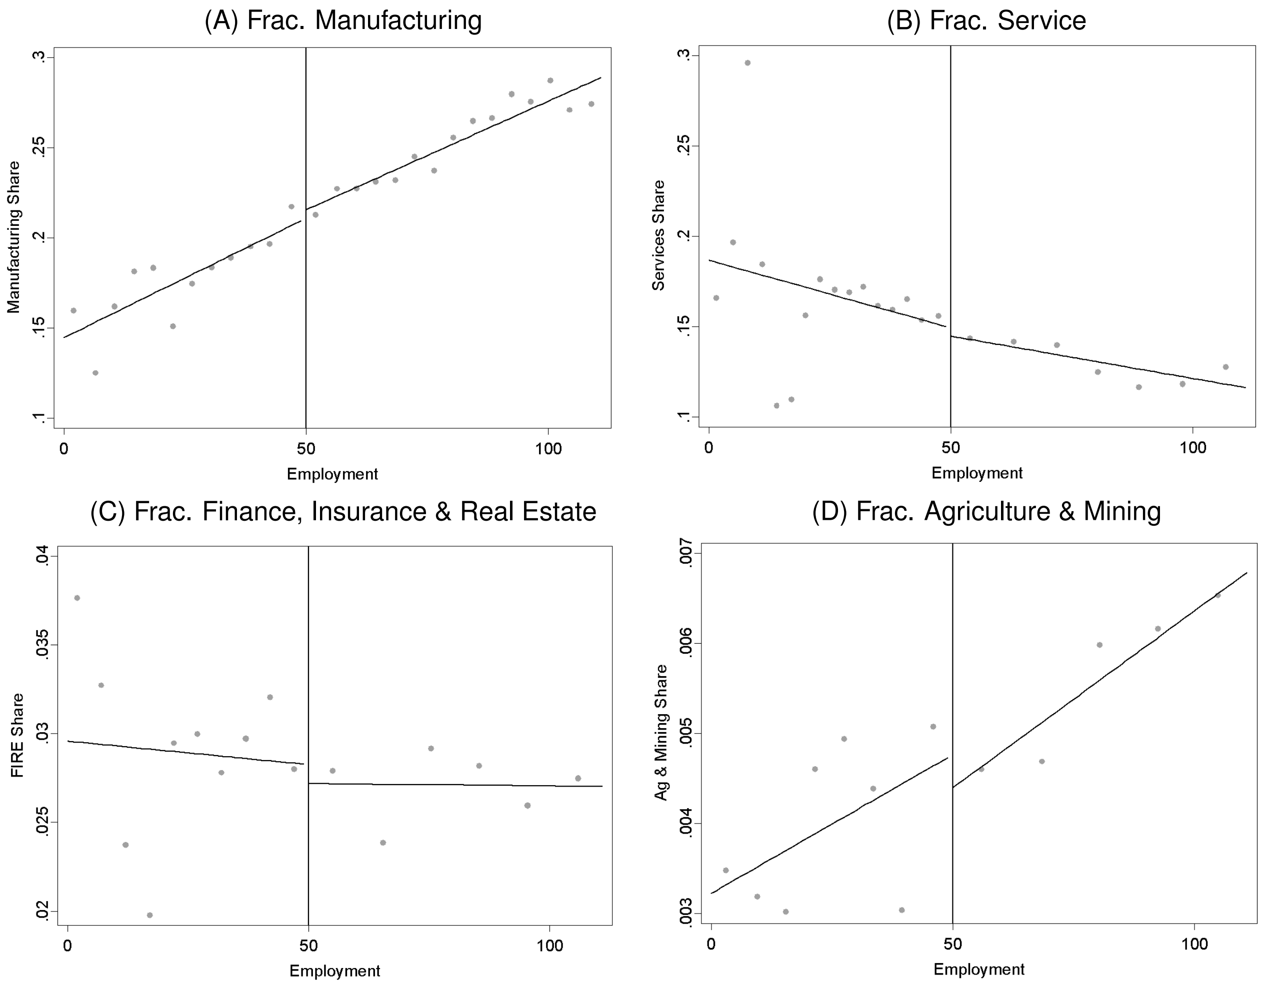

Notes: The figures illustrate the average fraction of manufacturing (panel A), service (panel B), finance, insurance, and real estate (panel C), and agriculture and mining firms (panel D) by firm employment. For these figures we choose the optimal bin size to the left and the right of the 50 employee cutoff based on the methods described in Calonico et al. (2015). The lines are fitted linear regressions. The graphs are based on the UK Quarterly Capital Expenditure Survey (QCES).

of productivity-enhancing investments that may have impacted workers. The fact that firms responded solely via increased ICT investment (as we show next), and not via other types of capital investment, seems to rule out an impact via the firms' "cash on hand".

In addition, we not only exploit the targeted nature of the tax incentive but also its timing. The incentive was introduced in 2000 and expired in 2004. If the effects that we identify are truly due to the tax incentive, then we would expect to find no visible discontinuity in outcomes prior to the introduction of the policy, and we present evidence on this throughout. While the expected effect after the expiration of the tax credit is

Figure 3: Non-ICT Investment & Firm Age: 2001-2004

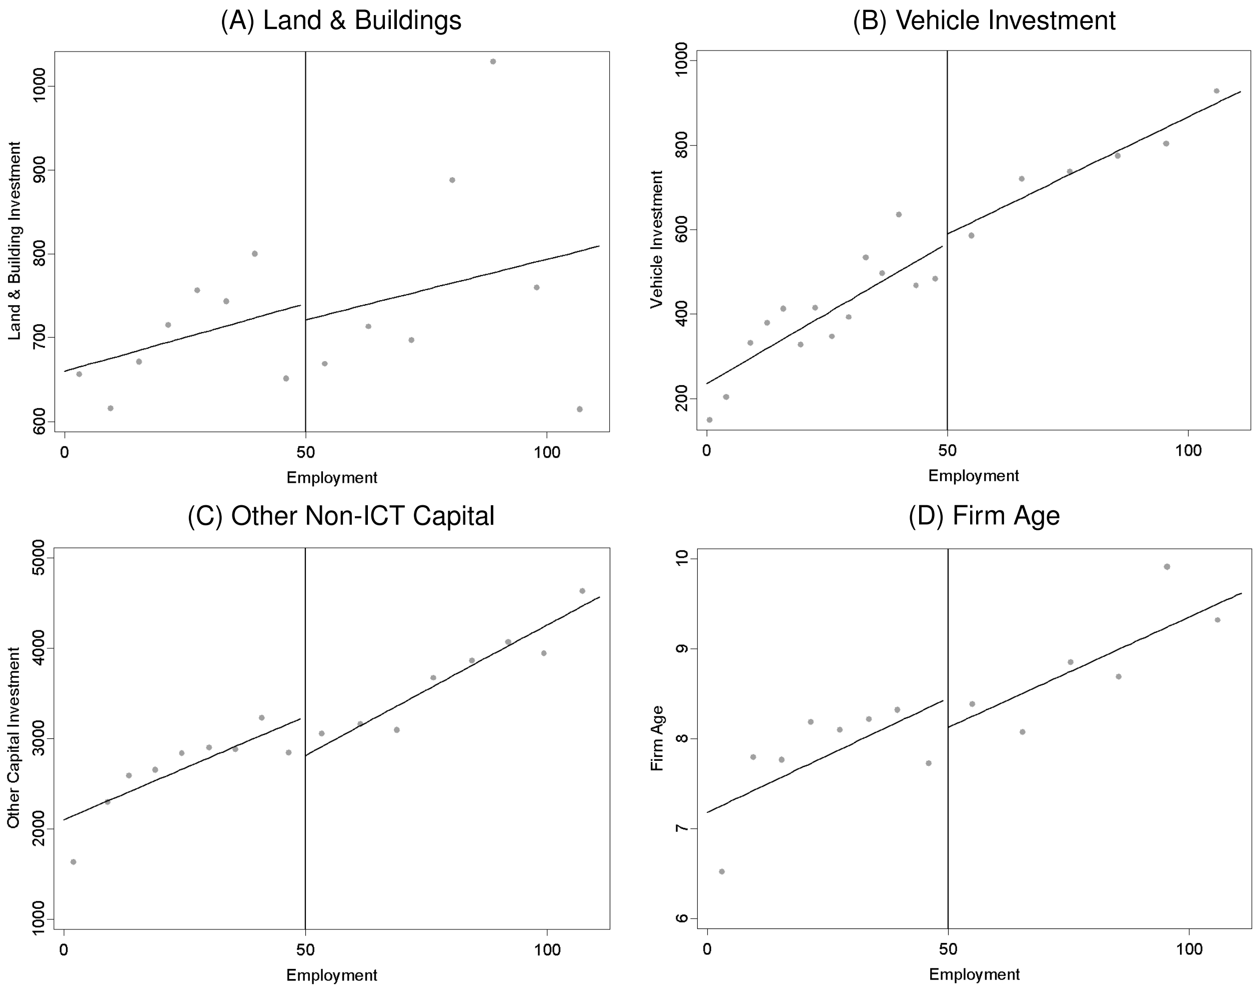

Notes: The figures illustrate investment in land and buildings (panel A), vehicles (panel B), and other non-ICT capital (panel C) by firm employment. Panel (D) illustrates firm age along the employment distribution. For these figures we choose the optimal bin size to the left and the right of the 50 employee cutoff based on the methods described in Calonico et al. (2015). The lines are fitted linear regressions. The graphs are based on the UK Quarterly Capital Expenditure Survey (QCES).

a-priori ambiguous for almost all outcomes, we will nevertheless estimate and discuss the observed effects.

Finally, we consider an additional approach in assessing the validity of our research design. We first note that the tax incentive only applied to private firms but not public firms. As a placebo test, we run our main analysis on public sector firms only, and find that there is no meaningful discontinuity in any of our key outcomes of interest at the 50 employee threshold, lending further credibility to our identification strategy. These results are reported in the Online Appendix.

Figure 4: Avg. Plant & Machinery Tax Deductions Claimed

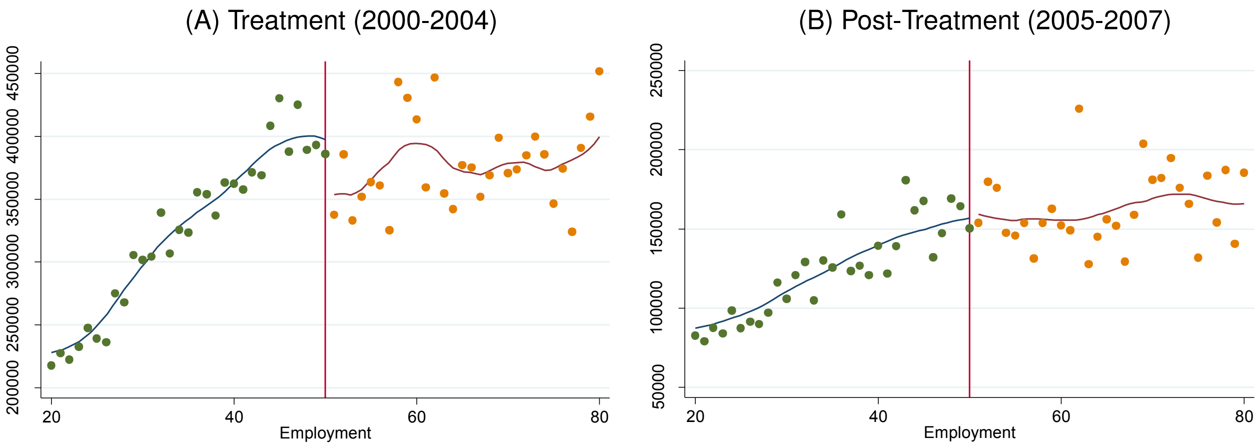

Notes: The figures plot average first-year tax allowances claimed by firm employment in British Pounds (GBP). Panel A pools the treatment period (2000-2004) while panel B pools the three years after the ICT tax incentive expired (2005-2007). The solid lines are locally weighted regressions.

5. Results

This section presents our main empirical results, starting with "first stage" estimates of the tax incentive's effect on firm-level ICT investment in Section 5.1, followed by "second stage" estimates of the effect on the employment and wage distribution in Section 5.2. Finally, Section 5.3 investigates the role of complementary organizational changes.

5.1. Firm Response to the Tax Incentive

We begin with graphical evidence from HMRC corporate tax returns over the period of interest. Figure 4 plots the average value of first year plant and machinery allowances claimed in a year by firm size, for the period 2000-2004 (treatment) as well as 2005-2007 (post-treatment).[23] Panel A of Figure 4 illustrates a sizable discontinuity in tax allowance claims at the 50 employee threshold. In contrast, panel B of Figure 4 indicates no notable discontinuity at the 50 employee threshold after the incentive expired at the end of 2004.[24] We view this as particularly convincing evidence of the success of the tax incentive, and of the

[23]Note that we treat each firm-year as an observation. Aggregating each firm's investment over the period produces nearly identical results. We did not implement the Calonico et al. (2015) procedures for the HMRC dataset due to HMRC constraints.

[24]The reader may notice that the *levels* of tax claims differ substantially between the periods. The primary reason for this is that several tax incentives targeting small and medium sized firms were available at various points during the 2000-2004 period, none of

Figure 5: ICT Investment by Firm-Level Employment

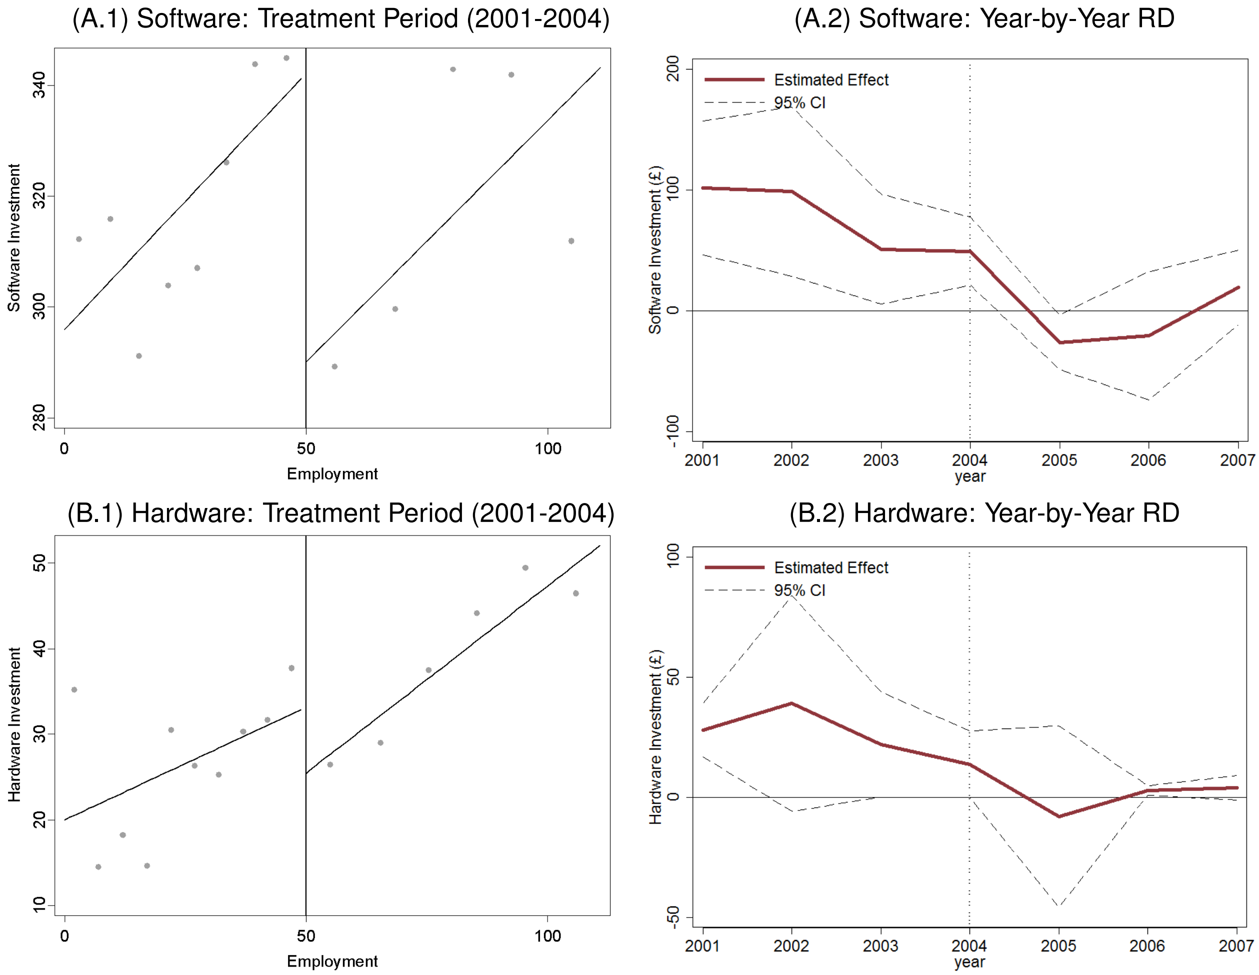

Notes: Panel A illustrates software investment along the firm size distribution for the Treatment period (2001-2004, panel A.1) as well as year-by-year local-linear RD estimates with associated confidence intervals for our QCES sample period (2001-2007, panel A.2). Panel B shows analogous figures for hardware investment. The optimal bin size is based on the methods described in Calonico et al. (2015). The lines in panels A.1 and B.1 are fitted linear regression lines. Optimal bandwidth selection, robust estimation, and inference for panels A.2 and B.2 are based on Calonico et al. (2014). The underlying data are from the UK Quarterly Capital Expenditure Survey (QCES).

validity of our research design. For instance, while firms' investment response to the tax incentive (estimated below) might be expected to spill over to some extent into the post-treatment period, the same cannot be said of firms' tax claims, and our finding of no effect in the post-treatment period is therefore reassuring.

which was available after 2004 (a full list and a history of U.K. capital allowances is available at https://www.gov.uk/hmrc-internal-manuals/capital-allowances-manual/ca10040). During the 2004-2007 period the UK had no investment incentive targeting small and medium sized firms specifically (though there were incentives available to all firms). In 2008 the capital allowance system was completely overhauled in favor of a less distortionary system.

While Figure 4 suggests that firms indeed responded to the tax incentive, at least in terms of their corporate tax filings, it is still possible that firms on either side of the threshold made similar investments in ICT over the period—i.e., it is possible that the incentive had no differential impact on actual ICT investment, but simply allowed those with 50 or fewer employees to claim allowances on investments they would have made even in the absence of the incentive. We therefore turn to evidence from the QCES, which reports quarterly investments in hardware and software by U.K. firms beginning in 2001. Panels A.1 and B.1 in Figure 5 plot the average investment in software and hardware against employment for the treatment period (2001-2004) while panels A.2 and B.2 plot the corresponding year-by-year RD estimates based on specification (3) for the entire QCES sample over 2001-2007.[25] These figures indicate a clear discontinuity in both software and hardware investment at 50 employees for the treatment period.[26] However, the year-by-year estimates indicate that the discontinuity was strongest immediately after the initial announcement of the policy, but slowly tapered off and completely vanished by 2005. We further note that the slopes of the linear regression fit to the left and right of the 50 employee cutoff are essentially identical for software investment but that there appears to be a small difference for hardware investment. However, this seems to be driven by unusually large reported average investment by the smallest firms, as can be seen in Panel B.1. Therefore, a local linear estimate that puts less weight on data points far from the 50 employee threshold appears appropriate, which is the primary specification we report throughout the text.

Panel A of Table 2 reports the corresponding local linear point estimates for the treatment period.[27] These estimates indicate that the ICT incentive on average induced approximately £71 thousand in additional annual software investment and about £13 thousand in additional annual hardware investment for just-eligible firms relative to just-not-eligible firms. For both types of ICT investments, these effects correspond to roughly one half of a standard deviation in investment for eligible firms.[28] Finally, we report

[25] Unfortunately, the QCES data are not available for the pre-2000 period (nor for 2000) so we cannot perform a falsification test during the pre-treatment period.

[26] Note that the optimal number of bins for these graphs was chosen according to the procedure from Calonico et al. (2015). In short, the procedure chooses the summary representation of the population regression function in the plots following a data-driven procedure. Furthermore, since the regression function to the left of the cutoff is allowed to differ from that to the right, the representation of the population regression function in the plots and the optimal bin size are chosen separately on either side of the size threshold.

[27] As mentioned above, throughout the paper we only report the point estimates for the local linear specification with the optimal bandwidth based on Calonico et al. (2014) in the main text. We experimented with a host of alternative specifications and all results presented here are robust to these modifications, which are reported in detail in the Online Appendix.

[28] From Table 1 we see that the the standard deviation for software and hardware investment for eligible firms is, 151.11 and

Table 2: Discontinuity in ICT Investment at 50 Employees

	A. Treatment: 2001 - 2004			B. Post Treatment: 2005-2007		
	(A.1) Softw.	(A.2) Hardw.	(A.3) Other	(B.1) Softw.	(B.2) Hardw.	(B.3) Other
RD	70.9158***	13.2712**	-14.5403	-27.6205	-10.442	-208.0195
	(24.2217)	(6.0726)	(233.4556)	(25.7599)	(9.5774)	(153.9913)
Obs.	75	81	107	75	123	81
Bandw.	37	40	56	37	72	40

Notes: The table reports local linear estimates of the regression coefficient τ_{RD} in model (3) based on the bias-corrected procedure described in Calonico et al. (2014). The dependent variables are (1) expenditures on software, (2) hardware, and (3) other capital investments. The optimal bandwidth is determined according to Calonico et al. (2014) and the data are aggregated to the optimal firm-size bin level. Robust standard errors based on Calonico et al. (2014) are reported in parentheses below each coefficient and significance levels are indicated by * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

a third category of ICT investment referred to as "other capital", which captures investments in non-ICT capital such as "fixtures/fittings, plant and machinery, etc.", according to the firm questionnaire. These investments were not eligible for the tax incentive and, reassuringly, we find no significant effect (see also panel C of Figure 3).

Panel B of Table 2 presents point estimates suggesting that, if anything, there was a relative underinvestment by small firms over the period 2005-2007. As noted above, it is not clear *a priori* whether we should expect the discontinuity in ICT investment to persist beyond the duration of the tax incentive, and the estimates in Table 2 suggest that in fact there was a subsequent fall in investment once the policy expired. This is consistent with firms' use of the tax incentive in order to manipulate the timing of planned investments—i.e., it is consistent with front-loading of planned ICT expenditure. This is an interesting finding from a policy perspective. To the extent that a tax incentive is designed to boost spending in "bad times", such a change in the timing of investments reflects precisely the intended effect of the policy. On the other hand, if the goal is to promote the use of ICT within small firms in order to foster long term productivity gains then a mere shift in the timing of pre-planned investments may be a less desirable outcome.

To illustrate this point more succinctly, panels A.2 and B.2 of Figure 5 document the evolution of the impact on software and hardware investment over the period 2001-2007, by plotting the RD estimates for each year separately, along with confidence intervals.[29] This exercise suggests that the primary impact was

23.61, respectively. Thus, the point estimates correspond to 0.47 and 0.56 standard deviations, respectively.

[29]Recall that we do not have QCES investment data for the pre-2001 period.

in the initial years of the policy while it tapered off and vanished by 2005. In fact, we note that there is some suggestion that investment declined in the immediate post-period, though the effect is only marginally significant in 2005.[30]

5.2. The Short-Run Effect of ICT on Labor Demand within the Firm

The finding of a positive effect of the tax incentive on ICT investment is perhaps not surprising, particularly considering that it took place during a period in which ICT was becoming increasingly important to day-to-day business operations in the U.K. In this section we go further and exploit the variation generated by these investments in order to estimate the short-run effect on hours worked and wages of workers employed in occupations with differing task content. Specifically, we use data from the ASHE and repeat the regression discontinuity approach from the previous sections with workers' weekly wages and weekly hours worked as the outcomes of interest.[31]

As in the analysis above we begin with a graphical depiction of the aggregate effect of the policy. Column 1 of Figure 6 plots the average weekly wage (panel A) and average weekly hours worked (panel B), as reported in the ASHE, against the running variable (firm size). These figures display clear discontinuities during the treatment period for both outcomes. To illustrate the evolution of the estimated effect over time, column 2 reports year-by-year RD estimates based on specification (3), alongside corresponding confidence intervals. As with the investment responses, this exercise highlights that these discontinuities are neither evident before the introduction nor after the expiration of the policy. Table 3 presents the corresponding local linear point estimates, pooled for the pre-treatment, treatment, and post-treatment periods, respectively, confirming the graphical evidence displayed in Figure 6.

Overall, these aggregate effects suggest that ICT investment induced additional employment and higher wages, on average, within the investing firms. One interpretation of these results is that there are productivity gains associated with the adoption of ICT technologies, which partly accrue to workers within these

[30]We note that this could be related to a "recession-like" experience around 2005 in the UK, documented by Carrillo-Tudela, Hobijn, She and Visschers (2016), which may have disproportionately affected smaller, rather than larger firms.

[31]It is worth noting that, since our running variable is firm employment, it might at first seem as if any effect of the incentive on hours worked could potentially bias our results. In other words, the tax incentive itself may lead to employment growth, which may push marginal firms (those with 50 employees, for instance) over the tax incentive threshold. However, it is important to note that firms can only apply for the tax incentive once (it is a First Year Allowance). As a result, to the extent that the tax incentive induces growth in firm size, this can have no future effect on firm eligibility.

Figure 6: Average Earnings, Hours, & Labor Productivity

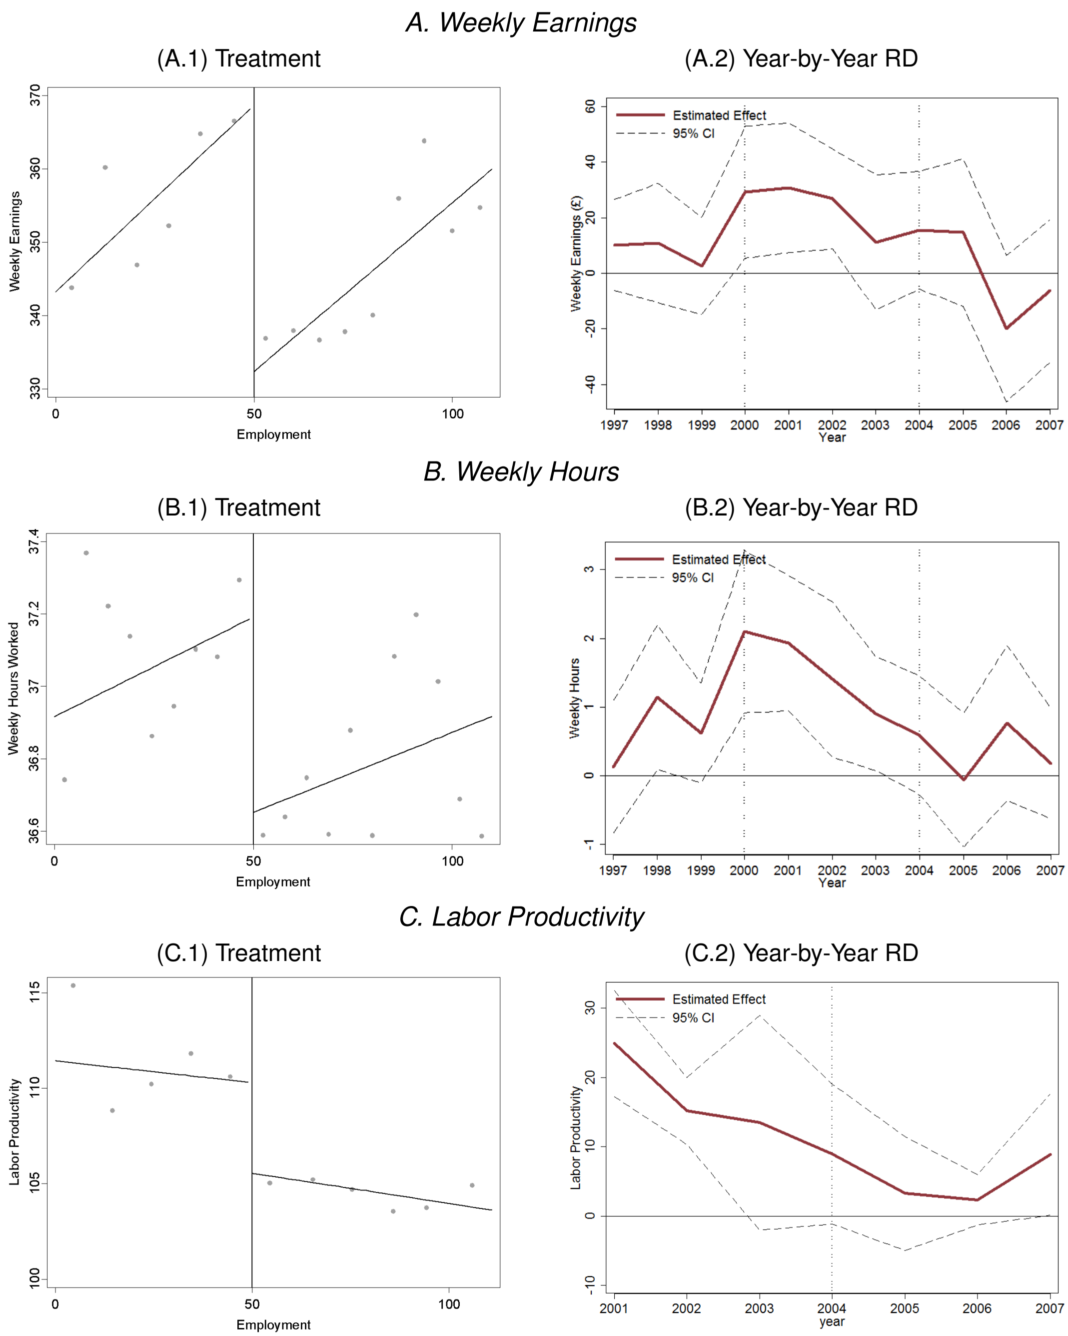

Notes: Column 1 displays linear RD plots during the treatment period (2001-2004) for weekly earnings (A.1), weekly hours (B.1) from the ASHE and labor productivity (revenue per worker, C.1) from the QCES. Column 2 shows corresponding year-by-year local linear RD estimates, for all available years in the corresponding data sources (ASHE and QCES). For figures A.1, B.1, and C.1 we chose the optimal bin size following Calonico et al. (2015). The solid lines in these figures are fitted linear regressions. Optimal bandwidth selection, robust estimation, and inference for panels A.2, B.2, and C.2 are based on Calonico et al. (2014).

Table 3: Discontinuity in Avg. Worker Outcomes

	A. Treatment: 2001 - 2004			B. Post Treatment: 2005-2008		
	(A.1) Earnings	(A.2) Hours	(A.3) Lab. Prod.	(B.1) Earnings	(B.2) Hours	(B.3) Lab. Prod.
RD	23.0448***	0.8879***	11.3953***	4.6277	0.1901	0.2638
	(4.5754)	(0.2576)	(3.4231)	(6.3052)	(0.2387)	(2.8218)
Obs.	101	95	67	112	81	67
Bandw.	50	47	33	61	40	33

Notes: The table reports local linear estimates of the regression coefficient τ_{RD} in model (3) based on the bias-correction procedure described in Calonico et al. (2014). The dependent variables are (1) weekly earnings, (2) weekly hours from the ASHE, and (3) and labor productivity (revenue/worker) from the QCES. The optimal bandwidth is determined according to Calonico et al. (2014) and the data are aggregated to the optimal firm-size bin level. Robust standard errors based on Calonico et al. (2014) are reported in parentheses below each coefficient and significance levels are indicated by * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

firms. Indeed, RD estimates with respect to QCES-reported labor productivity (revenue per worker) reveal a discontinuity during the treatment period but not after (see panel C of Figure 6 as well as columns A.3 and B.3 in Table 3).

In the context of the conceptual framework outlined in Section 3, the significant wage impact suggests that firms that took up the incentive likely faced an inelastic labor supply, which suggests that firms may have faced segmented labor markets. However, as noted in Section 3, one should keep in mind that there may be (mitigated) wage effects even when labor markets are fully integrated. In addition, the effect may mask underlying compositional changes in the demand for different types of work, such that the wage rise may not have been equally shared across the workforce. We explore these distributional consequences to the extent that our data allow in the next section.

Interestingly, the year-by-year RD estimates displayed in column 2 of Figure 6 suggest that the impact on workers observed in the treatment period was *not* sustained following the end of the policy. Given the durable nature of computer hardware and software, this is a surprising result and deserves further discussion. While we cannot be certain why the differential effects disappear, we note that the effect is estimated as the difference in worker outcomes relative to the control group, so that the finding may arise due to a quick convergence in pay and employment on the part of medium- and large-sized firms to the levels set by the incentive-impacted small firms. Since this was a period in which ICT-competent workers saw high demand and rising wages, and labor market turnover was at historically high levels (see Macaulay, 2003), labor market competition may have quickly eroded the differential effects between groups.

Figure 7: ICT Adoption & Timing of Innovation (2002-2004)

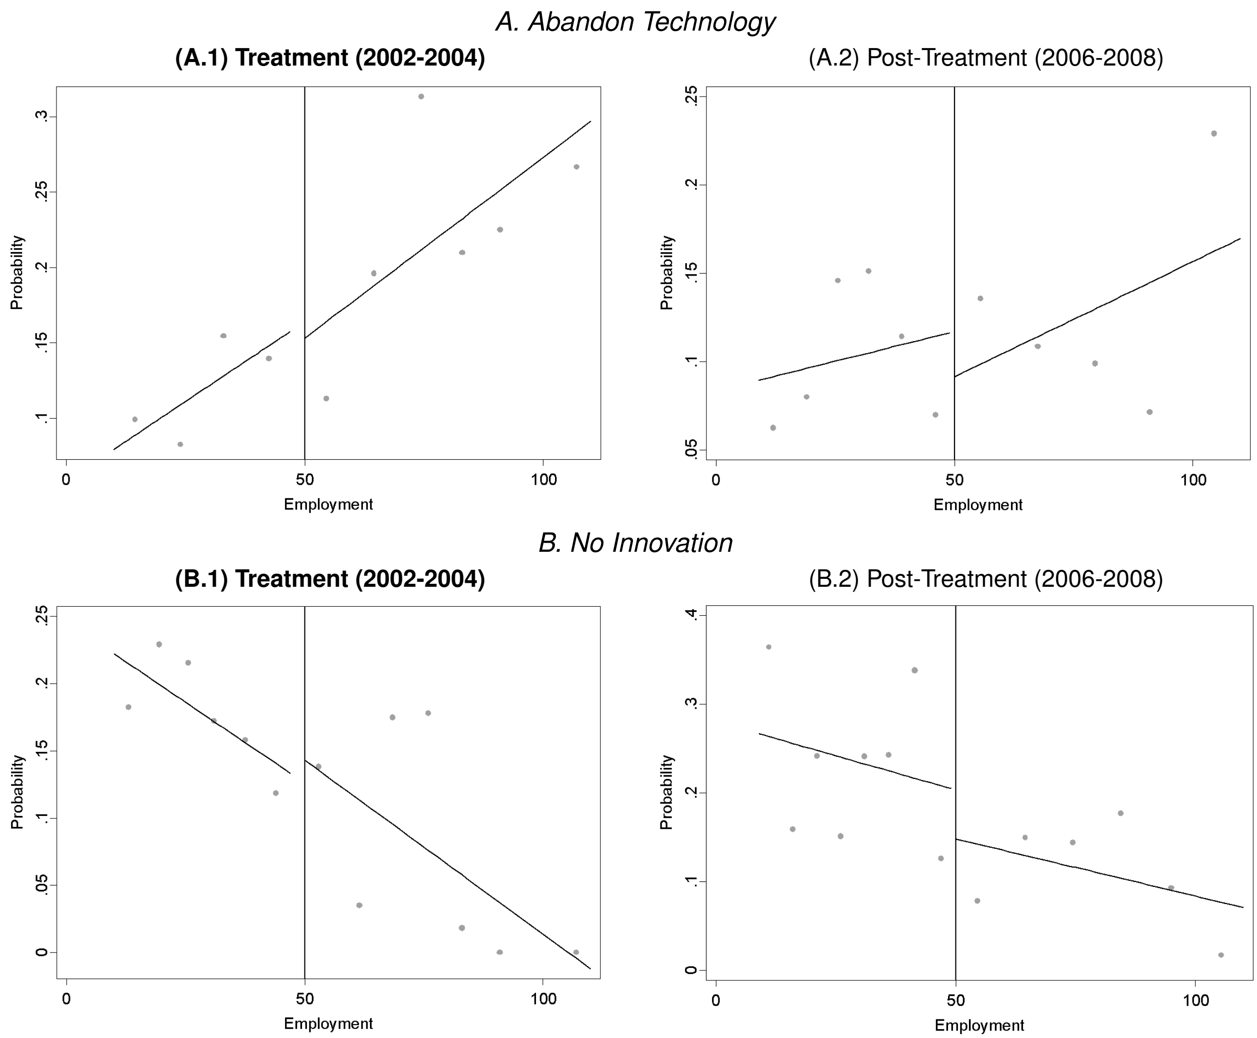

Notes: The figures plot the fraction of firms reporteing to have "abandoned innovations or technolgy" (panel A) or stated that they had "not made any new invstments in innovation" (panel B) reported in the CIS against firm size. Panels A.1 and B.1 present estimates for the series of questions pertaining to the period 2002-2004 (treatment) while panels A.2 and B.2 cover the period 2006-2008 (post-treatment). For all figures we chose the optimal bin size following Calonico et al. (2015). The lines are fitted linear regressions.

An alternative story is that the wage and employment impacts arose in part due to process innovations induced by the incentive that were subsequently abandoned (in Section 5.3 we show that firms indeed undertook organizational changes in response to the incentive). With this in mind, we can turn to the U.K. Innovation Survey (CIS) which, among other things, asks firms whether innovation activities were abandoned during a given period. Indeed, Figure 7 provides suggestive evidence for this possibility, since small firms appear to be more likely to have abandoned technologies or chosen not to innovate in the period after the tax incentive expired. It therefore seems that a subset of the firms who were induced to enact organizational changes in response to the incentive later abandoned those changes. However, ultimately, our findings represent something of a mystery and we leave further exploration to future research.

5.2.1. The Impact of ICT Investments on Workplace Tasks

We devote the remainder of this section to investigating the extent to which the productivity gains we found accrued to workers performing different workplace tasks. To this end, we estimate the effect of the policy on workers within the four broad occupation groups defined in Acemoglu and Autor (2011): (1) managerial, professional and technical occupations; (2) sales, clerical and administrative support; (3) production, craft, repair and operative occupations; and (4) service occupations. As these authors argue, the categories are broadly representative of different sets of production tasks, namely (1) non-routine, cognitive tasks; (2) routine, cognitive tasks; (3) routine manual tasks; and (4) non-routine manual tasks, respectively. As noted in Section 1, several studies support the idea that this classification of occupations reflects important dimensions of computer-worker interaction.

Our estimates indicate that the overall gains in both hours worked and wages (Figure 6), were by no means evenly shared across groups of workers performing fundamentally different tasks. In particular, Figures 8 and 9 along with Table 4 indicate that the largest impact of the tax incentive was on cognitive workers, with little evidence of an impact on manual workers. The latter finding seems plausible, given that manual work is not likely to be strongly complementary to, nor substitutable with, ICT. Taken together, and if we assume that the relative supply of task types over this period was fixed, these results imply a significant outward shift in the demand for non-routine, cognitive tasks and a simultaneous, but less pronounced, inward shift in the demand for routine, cognitive tasks. In the context of the conceptual framework outlined in Section 3, the observed wage impacts once again imply either that firms face an inelastic labor supply, and possibly segmented labor markets, or else that there are changes in the composition of workers within task groups that are unobservable to us; for instance, shifts toward occupations that are more (for the non-routine, cognitive group) or less (for the routine, cognitive group) complementary with ICT.[32]

We note that the slopes of the fitted regression lines in Figures 8 and 9 are largely identical on either side

[32]In theory we could further decompose these task groups into individual occupations, or some other disaggregation, but the sample sizes become quite small in many cases, and the estimates therefore less reliable.

Figure 8: Weekly Earnings: Routine, Non-Routine, Cognitive & Manual Jobs (2000-2004)

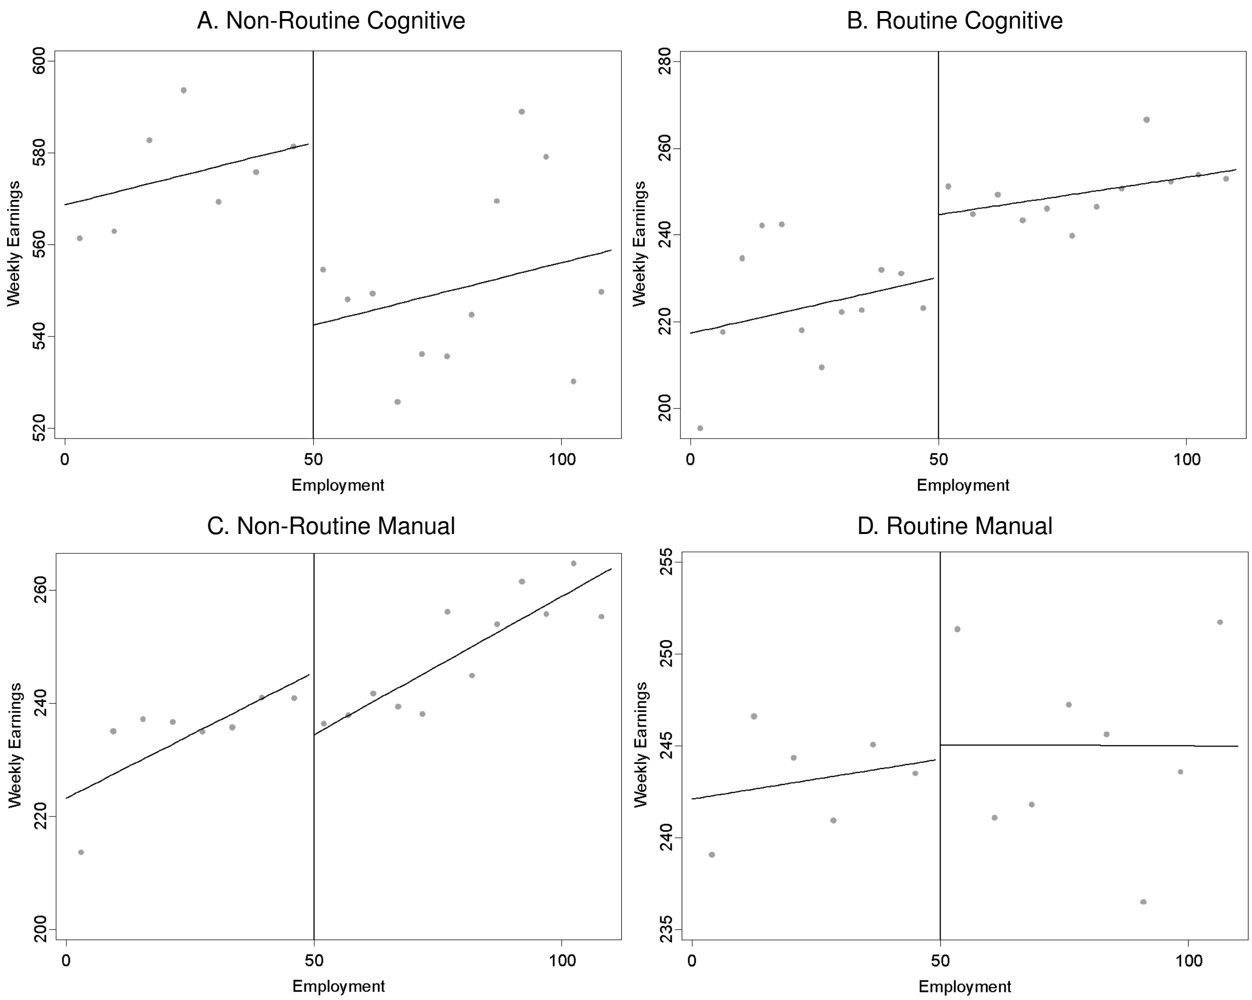

Notes: The figures decompose the aggregate results on weekly earnings from panel (A.1) of Figure 6 into the four occupation groups defined in Acemoglu and Autor (2011): non-routine cognitive (A), routine cognitive (B), non-routine manual (C), and routine manual (D). For all figures we chose the optimal bin size following Calonico et al. (2015). The lines are fitted linear regressions.

of the threshold for all specifications except hours worked within cognitive occupations (panels A and B of Figure 9). This raises the concern that a linear specification that exploits the entire range of data depicted in the Figures may lead us to falsely interpret non-linearities as discontinuities. To address this concern, all estimates presented in Table 4 are *local* linear estimates, with the optimal bandwidth around the policy cutoff chosen according to Calonico et al. (2014). In the Online Appendix, we report specifications in which we vary the bandwidth around the Calonico et al. (2014) optimum and we further estimate local second order polynomials instead of linear specifications. As an alternative graphical approach, we also construct

Figure 9: Weekly Hours: Routine, Non-Routine, Cognitive & Manual Jobs (2000-2004)

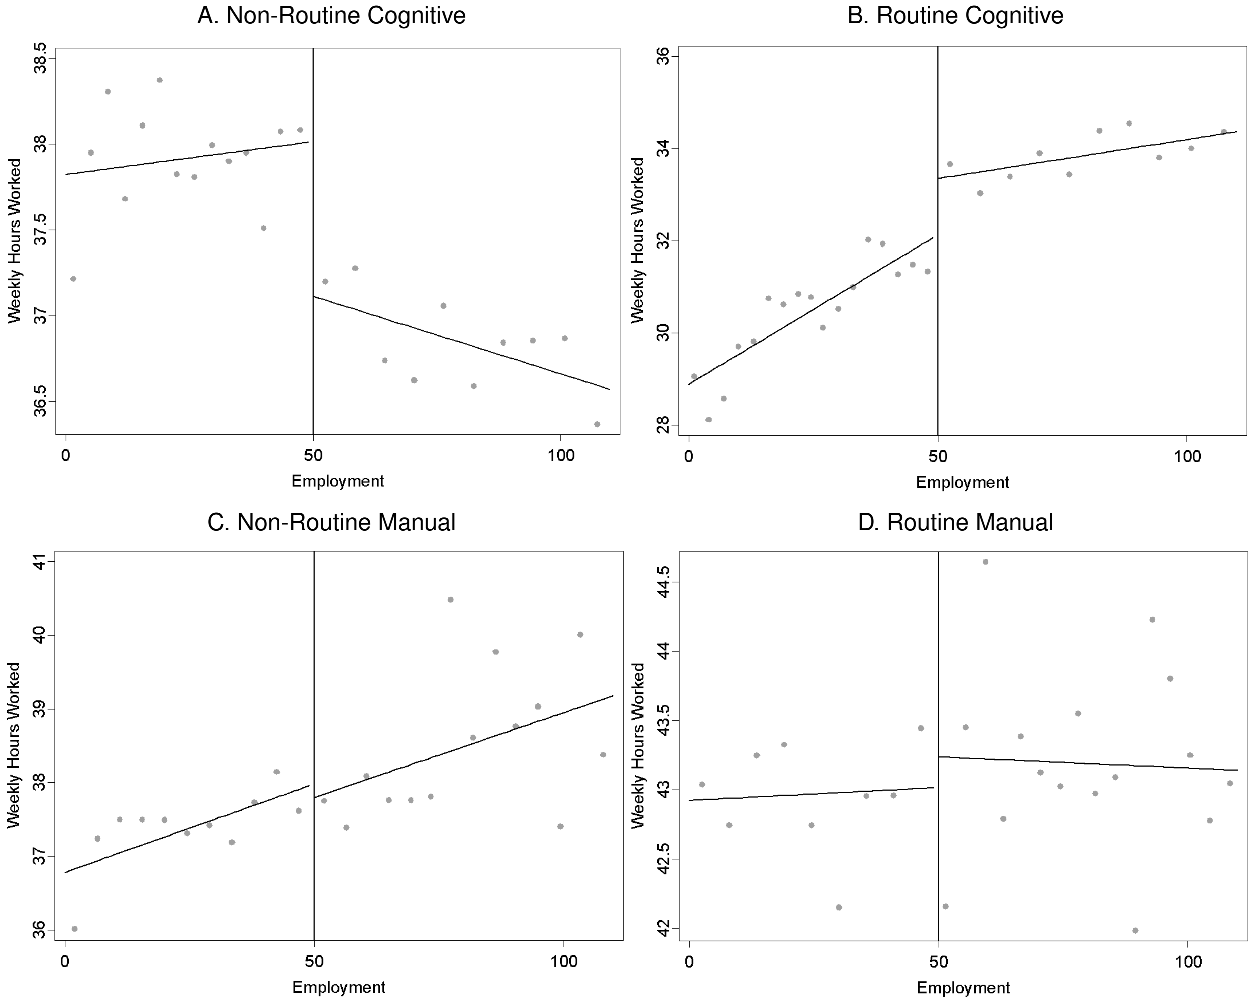

Notes: The figures decompose the aggregate results on hours from panel (B.1) of Figure 6 into the four occupation groups defined in Acemoglu and Autor (2011): non-routine cognitive (A), routine cognitive (B), non-routine manual (C), and routine manual (D). For all figures we chose the optimal bin size following Calonico et al. (2015). The lines are fitted linear regressions.

Figures analogous to those displayed here with fitted second order polynomials instead of linear regression lines. Reassuringly, all the results presented in the main text are robust to the menu of sensitivity analyses just described.[33]

Our estimates suggest that the earnings gains accruing to non-routine, cognitive workers are about five times as large as the losses for routine, cognitive workers, consistent with the observed net positive effect

[33] In fact, we conduct this same set of robustness checks for all main analyses reported in the text and these results are presented in the Online Appendix.

Table 4: Discontinuity in Worker Outcomes: Routine, Non-Routine, Cognitive & Manual

	A. Pre-Treat.: 1997 - 1999			B. Treatment: 2001 - 2004			C. Post Treat.: 2005-2008		
	(A.1) Earn-ings	(A.2) Hours	(A.3) Wage Disp.	(B.1) Earn-ings	(B.2) Hours	(B.3) Wage Disp.	(C.1) Earn-ings	(C.2) Hours	(C.3) Wage Disp.
A. Non-Routine Cognitive									
RD Estimate	10.61	0.46	0.51	**42.3821*** **	**0.5285** **	**0.341** **	8.9	0.14	0.83
Std. Err.	(10.4)	(0.35)	(0.91)	(11.6154)	(0.2215)	(0.1611)	(14.1)	(0.29)	(0.69)
Obs.	110	109	95	114	117	114	115	106	117
Bandwidth	59	58	47	63	66	63	64	55	66
B. Routine Cognitive									
RD Estimate	1.33	-0.38	0.28	**-8.6142***	**-0.5259***	**-0.2412***	4.02	0.27	0.44
Std. Err.	(4.9)	(0.49)	(0.34)	(4.797)	(0.307)	(0.1693)	(6.99)	(0.34)	(0.45)
Obs.	111	112	112	95	93	95	113	95	95
Bandwidth	60	61	61	47	46	47	62	47	47
C. Non-Routine Manual									
RD Estimate	-5.35	-0.49	-0.62	**4.5329**	**0.247**	**-0.4123**	8.12	-0.45	-0.57
Std. Err.	(4.16)	(0.49)	(0.73)	(4.3347)	(0.3536)	(0.3834)	(6.38)	(0.55)	(0.56)
Obs.	128	114	111	110	107	112	436	332	114
Bandwidth	77	63	60	59	56	61	58	41	63
D. Routine Manual									
RD Estimate	-5.99	0.12	-0.35	**-6.3721**	**-0.1107**	**-0.5343**	2.97	-0.17	-0.7
Std. Err.	(3.81)	(0.16)	(0.48)	(4.7223)	(0.5353)	(0.4809)	(4.92)	(0.53)	(0.84)
Obs.	115	120	105	116	111	105	116	105	116
Bandwidth	64	69	54	65	60	54	65	54	65

Notes: The table reports local linear estimates of the regression coefficient τ_{RD} in model (3) based on the bias-correction procedure described in Calonico et al. (2014). The dependent variables are (1) weekly earnings, (2) weekly hours, and (3) wage dispersion measured as the log wage gap between the 90th and 10th percentile of the earnings distribution within each group. The optimal bandwidth is determined according to Calonico et al. (2014) and the data are aggregated to the optimal firm-size bin level. Robust standard errors based on Calonico et al. (2014) are reported in parentheses below each coefficient and significance levels are indicated by * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

on average earnings within the firm (see Table 3). Similarly, the policy-induced ICT investment lead to an approximate 30-minute increase in the average weekly hours for non-routine, cognitive workers and about an equal decrease in hours worked by routine, cognitive workers. Note, however, that these average results do not allow us to disentangle hiring and firing decisions from the reallocation of existing workers to new tasks within the firm. That is, these average effects on hours and wages may result either from hiring non-routine, cognitive workers and firing routine, cognitive workers, or else from promoting a formerly routine, cognitive worker to a new job in which she performs primarily non-routine, cognitive tasks; for instance,

Figure 10: Weekly Earnings: Routine, Non-Routine, Cognitive & Manual Jobs: Year-by-Year RD

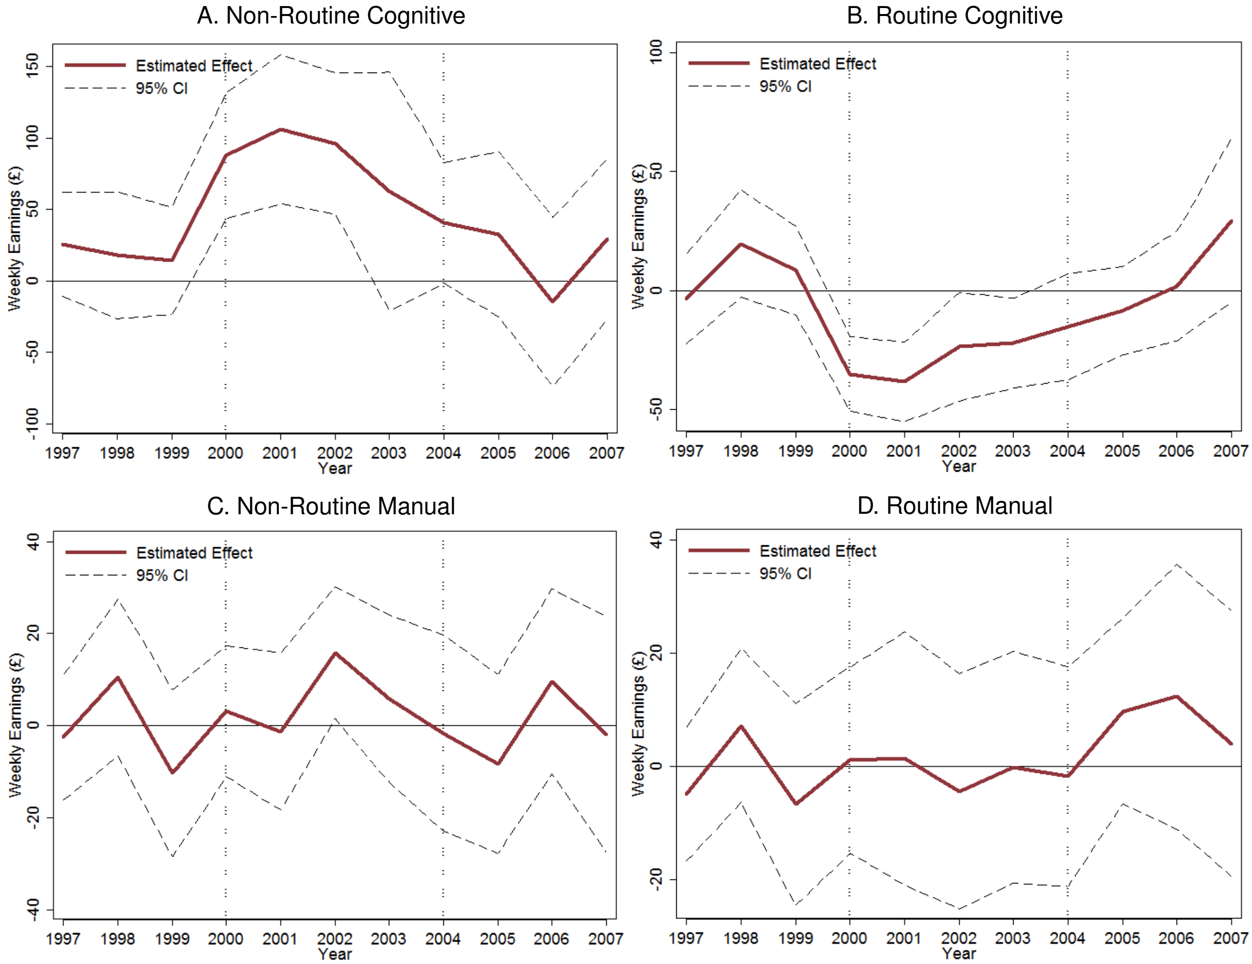

Notes: The figures decompose the year-by-year earnings results in from panel (A.2) of Figure 6 into the four occupation groups defined in Acemoglu and Autor (2011): non-routine cognitive (A), routine cognitive (B), non-routine manual (C), and routine manual (D). The figures report year-by-year local linear RD estimates and associated 95% confidence intervals. Optimal bandwidth selection, robust estimation, and inference are based on Calonico et al. (2014). The dashed vertical lines indicate the first and last year of the tax incentive.

giving computers and new job titles to all the secretaries in the office.

We again investigate the timing of these effects by plotting year-by-year local linear RD estimates (with associated confidence intervals) in Figures 10 and 11. Consistent with the pooled estimates reported in Table 4, we see that the wage and employment impacts on non-routine, cognitive workers and routine, cognitive workers are concentrated early in the policy period, and then progressively fade. As expected, there is no effect on other worker types.

Given the aggregate nature of these occupation groups, we next explore the impact of the policy on wage

Figure 11: Weekly Hours: Routine, Non-Routine, Cognitive & Manual Jobs: Year-by-Year RD

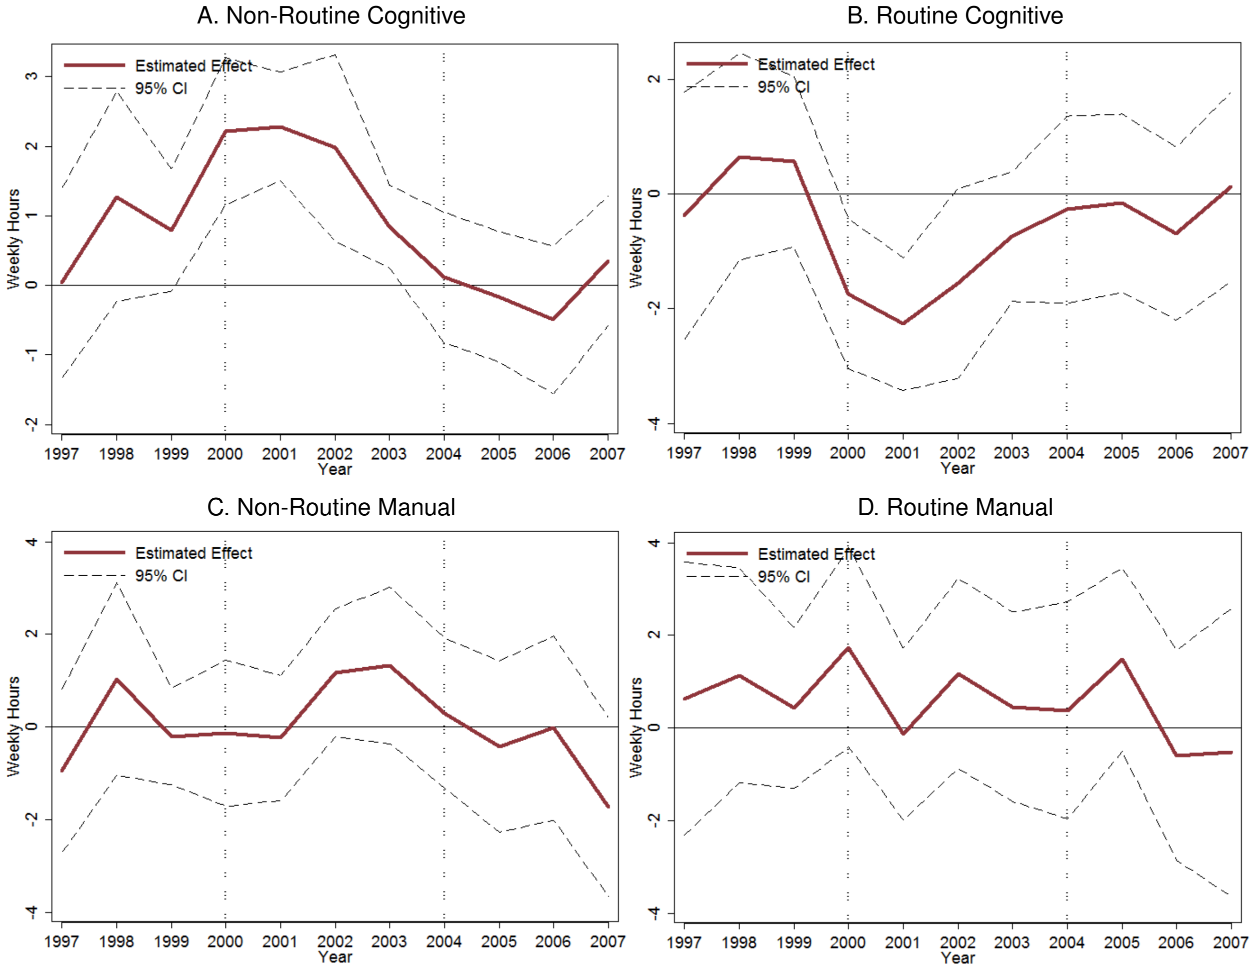

Notes: The figures decompose the aggregate year-by-year results on hours from panel (B.2) of Figure 6 into the four occupation groups defined in Acemoglu and Autor (2011): non-routine cognitive (A), routine cognitive (B), non-routine manual (C), and routine manual (D). The figures report year-by-year local linear RD estimates and associated 95% confidence intervals. Optimal bandwidth selection, robust estimation, and inference are based on Calonico et al. (2014). The dashed vertical lines indicate the first and last year of the tax incentive.

dispersion within each occupation group. Specifically, we focus on the difference between the log wage at the 90th percentile of the wage distribution and the log wage at the 10th percentile (the so-called 90-10 gap) within each of the four occupation groups. Column 3 in each panel of Table 4 reports the estimated second stage effects while the associated year-by-year estimates are displayed in Figure 12. The results suggest an overall rise in wage inequality for non-routine, cognitive work as well as a mild decline in wage inequality for routine, cognitive work that is only borderline statistically significant. There is no evident effect within manual occupations. Thus, the ICT investments not only increased the average return to non-

Figure 12: Wage Dispersion: Routine, Non-Routine, Cognitive & Manual Jobs: Year-by-Year RD

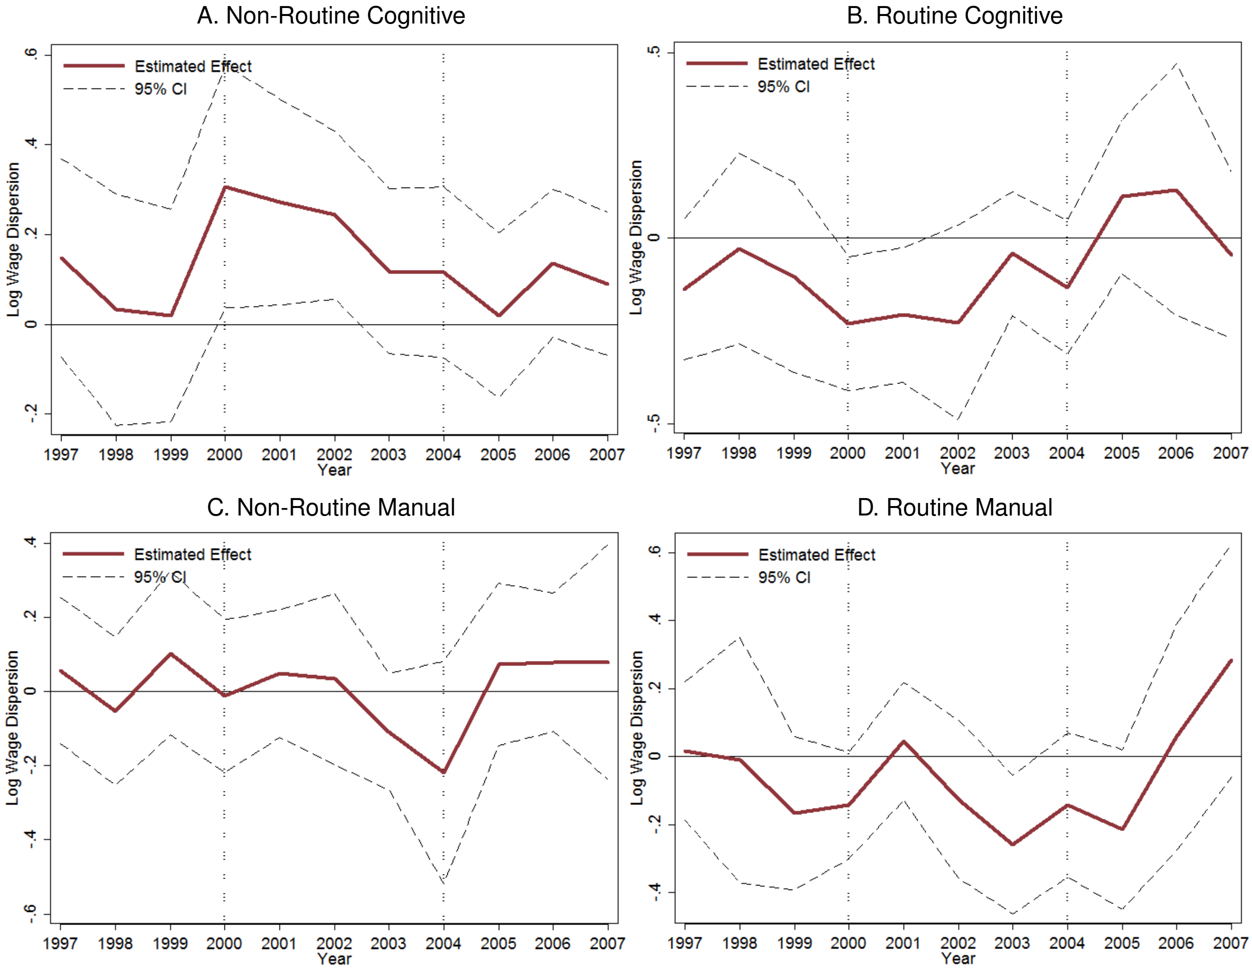

Notes: The figures report year-by-year RD estimates and associated 95% confidence intervals for wage dispersion (the log 90-10 percentile difference in weekly earnings) within the four occupation groups defined in Acemoglu and Autor (2011): non-routine cognitive (A), routine cognitive (B), non-routine manual (C), and routine manual (D). Optimal bandwidth selection, robust estimation, and inference are based on Calonico et al. (2014). The dashed vertical lines indicate the first and last year of the tax incentive.

routine, cognitive work but also increased the gap between the most and least productive workers. Again, Figure 12 highlights the now familiar pattern in which all significant effects are concentrated in the early part of the policy period and gradually fade. We will return to these results in Section 6 in order to compare two competing models for the effects of ICT adoption on workers.

5.3. Organizational Response to ICT Adoption

One potential explanation for the seemingly large effects on worker outcomes and labor productivity is the possibility of concurrent organizational change. While it is thought to be an important complement to the ICT-adoption process there is little evidence on the relative timing of organizational changes and technology adoption. For instance, it may be that firms make singular investments in new technologies and then integrate those technologies over a period of time in conjunction with incremental changes to the organization. Alternatively, organizational changes may be planned explicitly in order to coincide with ICT investments, and may therefore be enacted over a short horizon. In this case, a fall in the effective price of ICT may induce firms to immediately restructure their workforce in favor of ICT-complementary workers, both directly through hiring and firing as well as via other types of restructuring.

In this section we address the magnitude and timing of organizational changes with respect to ICT investment. Specifically, we present regression discontinuity results for firm responses to the following question in the U.K. Innovation Survey: "Did your enterprise make major changes in the following areas of business structure and practices during the three-year period, 2002-2004?" There are four aspects of business strategy addressed, for which firms could simply answer yes or no. These are: "New or significantly changed corporate strategy", "Implementation of advanced management techniques", "Implementation of major changes to your organizational structure", and "Implementation of changes in marketing concepts or strategies". We report summary statistics for the fraction of firms responding yes in panel D of Table 1.

We repeat the RD strategy from the previous sections, where now the variable of interest is the fraction of firms who responded "yes" to each question. Figure 13 and the point estimates in Table 5 suggest that the tax policy indeed led firms to implement advanced management techniques and to otherwise make changes to their organizational structure, while there is little evidence for the other channels. This strongly suggests that ICT adoption, changes in organizational structure, and changes in management practices are indeed both contemporaneous and complementary, a result that perhaps goes some way toward explaining the observed wage and employment effects. Consistent with our results in the previous sections, there are no notable effects in the post-treatment (2006-2008) period.

We also note that the survey question regarding organizational changes prompts survey respondents to consider offshoring of production activities as a potential activity that should lead to a "yes" response. Our

Figure 13: ICT Adoption & Organizational Change (2002-2004)

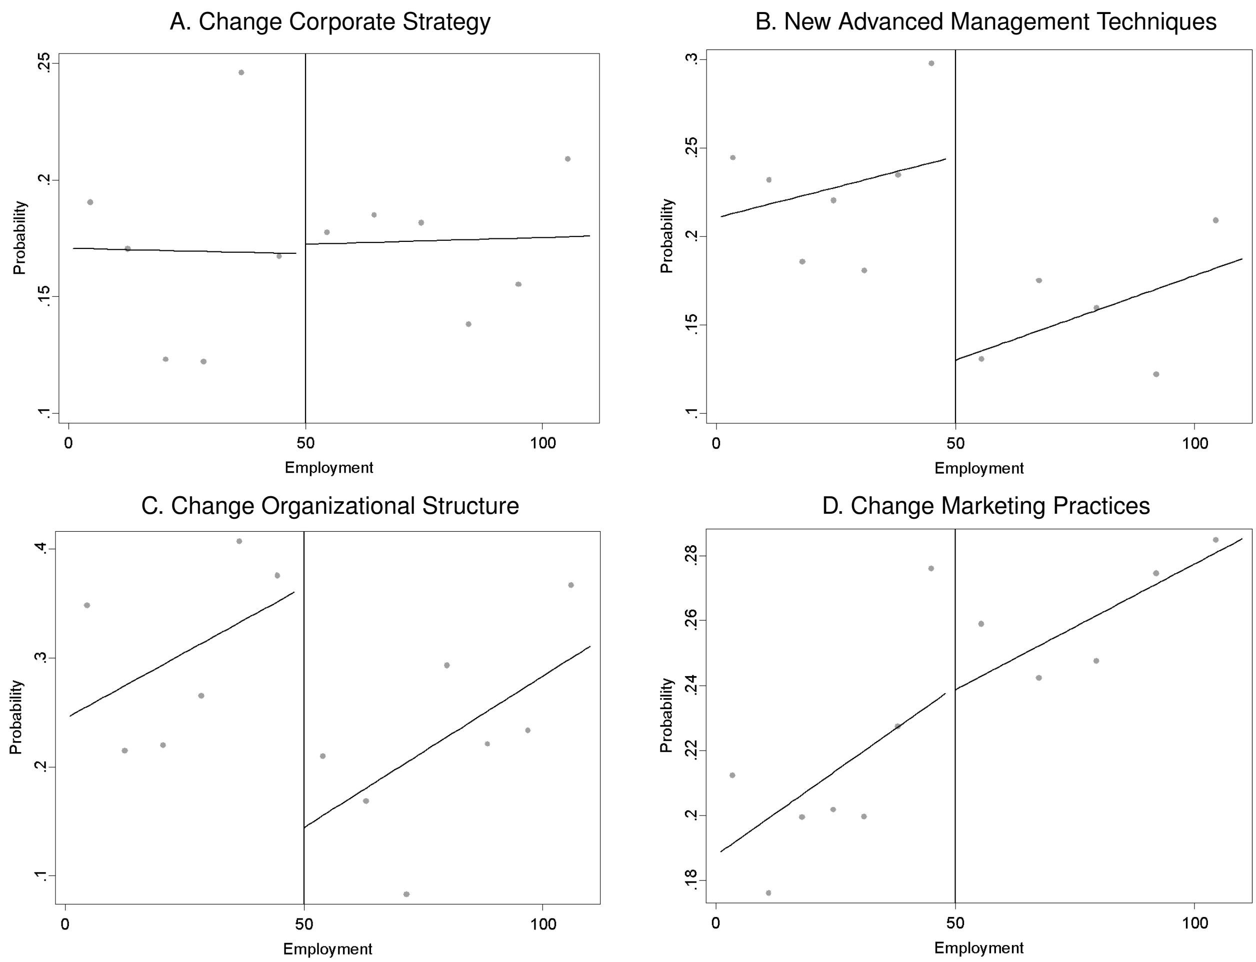

Notes: The figures report the fraction of firms how have conducted four types of organizational changes during 2002-2004 based on the CIS: (A) adopted a new corporate strategy, (B) adopted new advanced management techniques, (C) undertook a change in the organizational structure, or (D) adopted new marketing practicies. The optimal bin size was chosen based on Calonico et al. (2015). The lines are fitted linear regressions.

finding of increased organizational change due to the incentive is therefore consistent with the finding by Abramovsky and Griffith (2006) that ICT facilitates offshoring (though it is of course in no way definitive). And finally, in light of the labor demand response described above, the results effectively confirm the non-routine, cognitive bias of organizational change.[34]

[34]Caroli and Reenen (2001) were perhaps the first to identify this relationship.

Table 5: ICT Adoption & Organizational Change

	A. Treatment: 2002 - 2004				B. Post Treatment: 2006-2008			
	(A.1) Corp. Strat.	(A.2) Man. Tech.	(A.3) Org. Struc.	(A.4) Market-ing	(B.1) Corp. Strat.	(B.2) Man. Tech.	(B.3) Org. Struc.	(B.4) Market-ing
RD Estimate	0.0316	0.0661**	0.1773***	-0.0307	-0.0386	-0.0157	0.0138	-0.0051
	0.0654	0.0331	0.0629	0.0714	0.0493	0.0549	0.0769	0.0529
Obs.	85	79	79	79	84	88	84	87
Bandw.	42	39	39	39	42	47	42	45

Notes: The table reports local linear estimates of the regression coefficient τ_{RD} in model (3) based on the bias-correction procedure described in Calonico et al. (2014). The data are taken from two waves of the CIS. The dependent variables are indicator variables equal to one if a firm reports to have (1) implemented a change to its corporate strategy, (2) adopted new advanced management techniques, (3) implemented changes in the organizational structure, or (4) changed its marketing practices. Panel A reports results for the series of questions covering the period 2002-2004 while panel B covers the period 2006-2008. The optimal bandwidth is determined according to Calonico et al. (2014) and the data are aggregated to the optimal firm-size bin level. Robust standard errors based on Calonico et al. (2014) are reported in parentheses below each coefficient and significance levels are indicated by * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

6. Evidence on Competing Models of ICT Adoption

The link between ICT investments and labor demand within the firm is illuminated by the finding of a role for organizational change. However, mapping the observed effects to an economic model with more structure than our conceptual framework is not straightforward. Recently, two models have been proposed to explain the potential impact of ICT adoption on firm organization and labor demand. First is the tasks-based model, a framework that was most recently extended by Acemoglu and Autor (2011) in order to endogenize the assignment of tasks to skill. For our purposes, the relevant parameter in this model is the productivity of capital in performing routine tasks. An increase in this parameter is therefore the analog of the fall in the effective price (unit cost) of ICT that we exploit in the empirics.

In the context of this model, our empirical results are consistent with complementarity between ICT and non-routine, cognitive tasks, such that new computer technologies simply augment the workplace productivity of the existing high-skill workers who perform these tasks. At the same time, routine, cognitive tasks may be straightforwardly replaced by ICT—for instance, many types of back-office work may be replaced by computer software. Furthermore, as the Acemoglu and Autor (2011) model highlights, the increased demand for non-routine, cognitive work may draw in relatively less-productive workers who formerly performed different tasks, and this may increase wage inequality among non-routine, cognitive workers. To the extent that these new workers formerly performed routine, cognitive tasks and were possibly the most productive of those workers, this will reduce wage inequality among routine, cognitive tasks, again consistent with the wage-dispersion results above.

In this narrative there is not necessarily a formal role for organizational change, though it is important to note that the firm may report organizational change due to the fact that tasks have presumably been redistributed within the firm. Thus, our wage findings from above, combined with a role for organizational change, may still be consistent with the tasks-based mechanism. In contrast, Garicano and Rossi-Hansberg (2006) present a model in which organizational change plays a central role, and where reductions in the cost of ICT alter the relative importance of managers versus workers within the firm. In that model, the effects due to reductions in the cost of communication technologies (the Communications component of ICT) are different than those due to reductions in the cost of knowledge-acquisition technologies (the Information component of ICT). In short, they show that increased use of communications technologies widens inequality among managers while reducing inequality among workers, and has an ambiguous effect on between-group (manager vs worker) inequality.[35] In contrast, increased use of information-acquisition technologies widens inequality within both groups of workers while also increasing between-group inequality. The parameter of interest in this case is the cost of asking others for help within the firm, which ICT reduces, and which maps intuitively to a fall in the effective price of ICT that we observe.

Our empirical findings are therefore potentially consistent with the former effect, in which ICT investments are oriented toward reducing communication costs within the firm, thereby increasing wage dispersion among managers—who perform non-routine, cognitive tasks—and decreasing wage dispersion among workers—who, if we set aside manual work since it is seemingly unaffected, perform routine, cognitive tasks. We also note that our empirical results indicated a rise in between-group inequality (via a rise in the return to non-routine, cognitive tasks and a fall in the return to routine, cognitive tasks), which is consistent with both channels in the Garicano and Rossi-Hansberg (2006) model.

It is important to note that the findings reported in both Acemoglu and Autor (2011) as well as Garicano and Rossi-Hansberg (2006) are derived from economy-wide equilibria. Since the period that we exploit in

[35]We should note that this is an overly simple characterization of their findings, in part because in the model a firm may have multiple layers of management.

Figure 14: Wage Dispersion: Managers

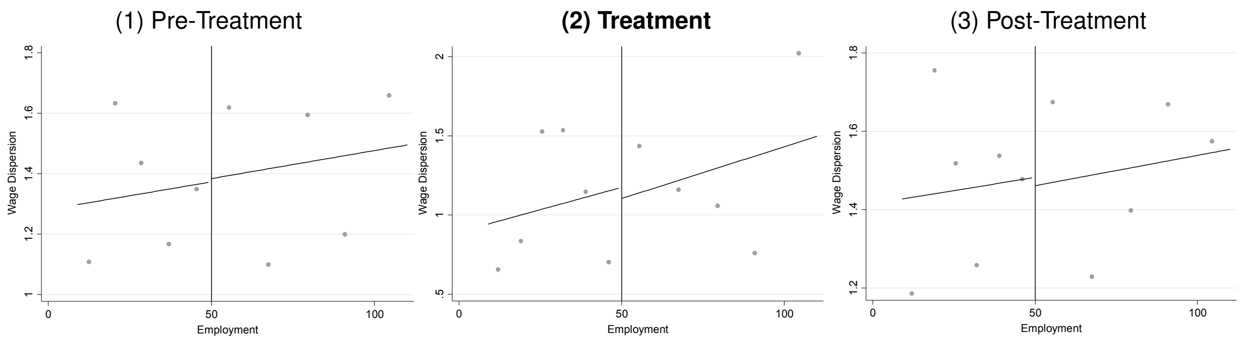

Notes: The figures plot the log wage gap between the 90th and 10th percentile of the earnings distribution within management occupations against frm size for three periods: pre-treatment (A, 1997-1999), treatment (B, 2000-2004), post-treatment (C, 2005-2007). The optimal bin size was chosen in accordance with Calonico et al. (2015). The lines are fitted linear regressions.

Table 6: Discontinuity in Worker Outcomes: Non-routine, cognitive & Managers

	A. Pre-Treat.: 1997 - 1999			B. Treatment: 2001 - 2004			C. Post Treat.: 2005-2008		
	(A.1) Earn-ings	(A.2) Hours	(A.3) Wage Disp.	(B.1) Earn-ings	(B.2) Hours	(B.3) Wage Disp.	(C.1) Earn-ings	(C.2) Hours	(C.3) Wage Disp.
A. Non-Routine Cognitive									
RD Estimate	10.61	0.46	0.51	**42.3821*** **	**0.5285** **	**0.341** **	8.9	0.14	0.83
Std. Err.	(10.4)	(0.35)	(0.91)	(11.6154)	(0.2215)	(0.1611)	(14.1)	(0.29)	(0.69)
Obs.	110	109	95	114	117	114	115	106	117
Bandwidth	59	58	47	63	66	63	64	55	66
B. Managers									
RD Estimate	10.64	0.46	0.45	**12.9022**	**0.1689**	**0.7445**	2.94	0.37	0.42
Std. Err.	(10.81)	(0.35)	(0.48)	(9.9151)	(0.2982)	(0.6872)	(4.92)	(0.53)	(0.34)
Obs.	110	109	95	116	105	116	116	105	116
Bandwidth	59	58	47	65	54	65	65	54	65

Notes: The table reports local linear estimates of the regression coefficient τ_{RD} in model (3) based on the bias-correction procedure described in Calonico et al. (2014). The dependent variables are (1) weekly earnings, (2) weekly hours, and (3) wage dispersion measured as the log wage gap between the 90th and 10th percentile of the earnings distribution within two groups of workers: panel A reports non-routine, cognitive workers; panel B reports management occupations. The optimal bandwidth is determined according to Calonico et al. (2014) and the data are aggregated to the optimal firm-size bin level. Robust standard errors based on Calonico et al. (2014) are reported in parentheses below each coefficient and significance levels are indicated by * $p < 0.1$, ** $p < 0.05$, and *** $p < 0.01$.

producing our estimates is relatively short, and we have no definitive evidence that labor markets are fully integrated, it is not necessarily the case that we are observing equilibrium outcomes. We therefore keep this in mind in interpreting our results.

In an effort to identify the relevant mechanisms, we exploit the narrow focus on the managerial class in

Garicano and Rossi-Hansberg (2006). Since this group is a subset of the non-routine, cognitive group, we can ask whether the wage patterns we observe for the non-routine, cognitive group are driven primarily by managers, as predicted by Garicano and Rossi-Hansberg (2006), or by other workers who perhaps use the technology more directly, which would be more consistent with the tasks-based framework. Specifically, we apply the RD strategy to the set of workers in the ASHE whose occupation codes classify them as managers, and report the results in Figure 14 and panel B of Table 6.

In contrast to non-routine, cognitive workers overall (panel A of Table 6), our estimates indicate no clear discontinuity in the managerial wage at the 50 employee threshold, suggesting that over the time period examined here the forces described by the knowledge hierarchy model may not be relevant. This points to a more direct augmentation and substitution of labor within the firm, consistent with the tasks model. Finally, we note that this exercise is clearly suggestive, and not definitive. Foremost, the mechanisms of the knowledge hierarchy model may be more dependent on general equilibrium forces relative to the tasks-based model. And, in addition, there may be important forces operating that are outside either of these models.

7. Concluding Remarks

Our short-run look at "what ICT does" refines the answer originally suggested by Autor et al. (2003). We find that the adoption of ICT leads to a rise in the demand for non-routine, cognitive tasks, even within a horizon of only five years. At the same time, we find a modest tendency for ICT to replace routine, cognitive work while manual work seems mostly unaffected.

We implement a research design that exploits exogenous variation in ICT investments generated by a temporary tax incentive. While economists have long sought to identify the conditions under which tax incentives are effective in stimulating investment demand (at least since Hall and Jorgenson's (1967) seminal work) a clear consensus has yet to emerge from this debate.[36] Most of the extant empirical work, even when exploiting natural experiments, takes a structural approach which often requires the approximation of many model quantities that are not directly observed in the data (e.g., firm- and asset-specific depreciation and tax rates, the rental rate of capital, etc.) and that are conditional on the particular structural assumptions. In

[36]For a few important contributions see Lucas (1976), Jorgenson and Yun (1990), Auerbach and Hassett (1991, 1992), Cummins et al. (1994, 1996), Goolsbee (1998), Chirinko et al. (1999), Cohen and Cummins (2006), as well as House and Shapiro (2008).

contrast, our research design relies on few identifying assumptions, which are likely to be satisfied.

The modest impact on routine workers reflected in our estimates suggests an asymmetry in the timing of the organizational change that goes along with ICT adoption. New technologies may demand immediate engagement by workers with the skill and ability to execute non-routine, cognitive tasks. As a result, organizational change—i.e., hiring and firing, extending worker hours, restructuring of workplace hierarchies, etc.—may align itself with this requirement in the short run. There is seemingly less of a need in the short run to replace workers who previously performed the routine tasks that are now performed by ICT—though, again, we do find evidence of some substitution. We also note that this asymmetry in timing is consistent with Jaimovich and Siu (2012), who find that about 92 percent of the routine jobs lost in the US since the 1980s were lost during a 12-month window following NBER dated recessions. This is despite the fact that aggregate investment is highly pro-cyclical, and therefore most investment—including ICT investment—happens during booms rather than immediate recoveries from a recession.

One of the standard predictions from neoclassical investment demand theory is that investment tax breaks—or, for that matter, any policy-induced reduction in the current price of investment goods—are only effective in altering the *timing* of investment. In this case, they will only stimulate current investment demand if they are expected to be sufficiently temporary and if the eligible assets are sufficiently long-lived. While the U.K. tax incentive studied here was indeed temporary, ICT capital is among the fastest depreciating forms of equipment, with annual depreciation rates of up to 30 percent. In light of this, the success of the ICT tax incentive explored here is somewhat puzzling. Some possible explanations include: (1) the tax incentive was quite generous (a 100 percent tax write-off); (2) computer purchases represent relatively small capital investments and therefore require less financial planning; (3) in 2000, the "computer revolution" was well under way and managers may have felt pressure to invest in ICT capital. We leave further exploration of this puzzle to future research.

Finally, we provide evidence on the relevance of recent theories of ICT adoption and wage inequality, finding that a tasks-based explanation best fits the patterns that we observe.

References

Abramovsky L, Griffith R. 2006. Outsourcing and Offshoring of Business Services: How Important is ICT? *Journal of the European Economic Association* **4**: 594–601.

URL https://ideas.repec.org/a/tpr/jeurec/v4y2006i2-3p594-601.html

Acemoglu D. 1998. Why Do New Technologies Complement Skills? Directed Technical Change And Wage Inequality. *The Quarterly Journal of Economics* **113**: 1055–1089.
URL http://ideas.repec.org/a/tpr/qjecon/v113y1998i4p1055-1089.html

Acemoglu D. 2002. Technical Change, Inequality, and the Labor Market. *Journal of Economic Literature* **40**: 7–72.
URL http://ideas.repec.org/a/aea/jeclit/v40y2002i1p7-72.html

Acemoglu D. 2007. Equilibrium Bias of Technology. *Econometrica* **75**: 1371–1409.
URL http://ideas.repec.org/a/ecm/emetrp/v75y2007i5p1371-1409.html

Acemoglu D, Aghion P, Lelarge C, Van Reenen J, Zilibotti F. 2006. Technology, information and the decentralization of the firm. Technical report, National Bureau of Economic Research.

Acemoglu D, Autor D. 2011. *Skills, Tasks and Technologies: Implications for Employment and Earnings*, volume 4 of *Handbook of Labor Economics*, chapter 12. Elsevier, 1043–1171.
URL http://ideas.repec.org/h/eee/labchp/5-12.html

Agha L. 2014. The effects of health information technology on the costs and quality of medical care. *Journal of health economics* **34**: 19–30.

Akcomak S, Kok S, Rojas-Romagosa H. 2013. The effects of technology and offshoring on changes in employment and task-content of occupations. CPB Discussion Paper 233, CPB Netherlands Bureau for Economic Policy Analysis.
URL http://ideas.repec.org/p/cpb/discus/233.html

Akerman A, Gaarder I, Mogstad M. 2015. The Skill Complementarity of Broadband Internet. *The Quarterly Journal of Economics* **130**: 1781–1824.
URL https://ideas.repec.org/a/oup/qjecon/v130y2015i4p1781-1824.html

Atasoy H. 2013. The effects of broadband internet expansion on labor market outcomes. *Industrial & Labor Relations Review* **66**: 315–345.

Auerbach AJ, Hassett K. 1991. Recent U.S. investment behavior and the tax reform act of 1986: A disaggregate view. *Carnegie-Rochester Conference Series on Public Policy* **35**: 185–215.
URL http://ideas.repec.org/a/eee/crcspp/v35y1991ip185-215.html

Auerbach AJ, Hassett K. 1992. Tax policy and business fixed investment in the United States. *Journal of Public Economics* **47**: 141–170.
URL http://ideas.repec.org/a/eee/pubeco/v47y1992i2p141-170.html

Autor DH, Dorn D. 2013. The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review* **103**: 1553–97.
URL http://ideas.repec.org/a/aea/aecrev/v103y2013i5p1553-97.html

Autor DH, Dorn D, Hanson GH. 2013a. The geography of trade and technology shocks in the united states. *American Economic Review* **103**: 220–25.
URL http://ideas.repec.org/a/aea/aecrev/v103y2013i3p220-25.html

Autor DH, Dorn D, Hanson GH. 2013b. Untangling trade and technology: Evidence from local labor markets. NBER Working Papers 18938, National Bureau of Economic Research, Inc.
URL http://ideas.repec.org/p/nbr/nberwo/18938.html

Autor DH, Levy F, Murnane RJ. 2003. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* **118**: 1279–1333.
URL http://ideas.repec.org/a/tpr/qjecon/v118y2003i4p1279-1333.html

Bartel A, Ichniowski C, Shaw K. 2007. How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *The Quarterly Journal of Economics* **122**: 1721–1758.
URL http://qje.oxfordjournals.org/content/122/4/1721.abstract

Beaudry P, Doms M, Lewis E. 2010. Should the personal computer be considered a technological revolution? evidence from u.s. metropolitan areas. *Journal of Political Economy* **118**: 988 – 1036.
URL http://ideas.repec.org/a/ucp/jpolec/doi10.1086-658371.html

Bertschek I, Cerquera D, Klein GJ. 2013. More bits–more bucks? measuring the impact of broadband internet on firm performance. *Information Economics and Policy* **25**: 190–203.

Black SE, Lynch LM. 2001. How To Compete: The Impact Of Workplace Practices And Information Technology On Productivity. *The Review of Economics and Statistics* **83**: 434–445.
URL https://ideas.repec.org/a/tpr/restat/v83y2001i3p434-445.html

Bøler EA, Moxnes A, Ulltveit-Moe KH. 2012. Technological Change, Trade in Intermediates and the Joint Impact on Productivity. CEPR Discussion Papers 8884, C.E.P.R. Discussion Papers.
URL https://ideas.repec.org/p/cpr/ceprdp/8884.html

Bresnahan TF, Brynjolfsson E, Hitt LM. 2002. Information Technology, Workplace Organization, And The Demand For Skilled Labor: Firm-Level Evidence. *The Quarterly Journal of Economics* **117**: 339–376.
URL http://ideas.repec.org/a/tpr/qjecon/v117y2002i1p339-376.html

Brynjolfsson E, Hitt LM. 2003. Computing productivity: Firm-level evidence. *The Review of Economics and Statistics* **85**: 793–808.
URL http://ideas.repec.org/a/tpr/restat/v85y2003i4p793-808.html

Calonico S, Cattaneo MD, Titiunik R. 2014. Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica* **82**: 2295–2326. ISSN 1468-0262.
URL http://dx.doi.org/10.3982/ECTA11757

Calonico S, Cattaneo MD, Titiunik R. 2015. Optimal data-driven regression discontinuity plots. forthcoming in the *Journal of the American Statistical Association* .

Cardona M, Kretschmer T, Strobel T. 2013. Ict and productivity: conclusions from the empirical literature. *Information Economics and Policy* **25**: 109–125.

Caroli E, Reenen JV. 2001. Skill-Biased Organizational Change? Evidence From A Panel Of British And French Establishments. *The Quarterly Journal of Economics* **116**: 1449–1492.
URL http://ideas.repec.org/a/tpr/qjecon/v116y2001i4p1449-1492.html

Carrillo-Tudela C, Hobijn B, She P, Visschers L. 2016. The extent and cyclicality of career changes: Evidence for the U.K. *European Economic Review* **84**: 18–41.
URL https://ideas.repec.org/a/eee/eecrev/v84y2016icp18-41.html

Chirinko RS, Fazzari SM, Meyer AP. 1999. How responsive is business capital formation to its user cost?: An exploration with micro data. *Journal of Public Economics* **74**: 53–80.
URL http://ideas.repec.org/a/eee/pubeco/v74y1999i1p53-80.html

Cohen D, Cummins J. 2006. A retrospective evaluation of the effects of temporary partial expensing. Finance and Economics Discussion Series 2006-19, Board of Governors of the Federal Reserve System (U.S.).
URL http://ideas.repec.org/p/fip/fedgfe/2006-19.html

Colombo MG, Delmastro M. 2004. Delegation of authority in business organizations: an empirical test. *The Journal of Industrial Economics* **52**: 53–80.

Cummins JG, Hassett KA, Hubbard RG. 1994. A Reconsideration of Investment Behavior Using Tax Reforms as Natural Experiments. *Brookings Papers on Economic Activity* **25**: 1–74.
URL http://ideas.repec.org/a/bin/bpeajo/v25y1994i1994-2p1-74.html

Cummins JG, Hassett KA, Hubbard RG. 1996. Tax reforms and investment: A cross-country comparison. *Journal of Public Economics* **62**: 237–273.
URL http://ideas.repec.org/a/eee/pubeco/v62y1996i1-2p237-273.html

De Stefano T, Kneller R, Timmis J. 2014. The (fuzzy) digital divide: The effect of broadband internet use on uk firm performance. Discussion papers, University of Nottingham, School of Economics.
URL http://EconPapers.repec.org/RePEc:not:notecp:14/06

DiNardo JE, Pischke JS. 1997. The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too? *The Quarterly Journal of Economics* **112**: 291–303.
URL http://ideas.repec.org/a/tpr/qjecon/v112y1997i1p291-303.html

Doms M, Dunne T, Troske KR. 1997. Workers, wages, and technology. *The Quarterly Journal of Economics* **112**: 253–90.
URL http://ideas.repec.org/a/tpr/qjecon/v112y1997i1p253-90.html

Doms M, Lewis E. 2006. Labor supply and personal computer adoption. Working Papers 06-10, Federal Reserve Bank of Philadelphia.
URL http://ideas.repec.org/p/fip/fedpwp/06-10.html

Draca M, Sadun R, van Reenen J. 2006. Productivity and ICT: A Review of the Evidence. CEP Discussion Papers dp0749, Centre for Economic Performance, LSE.
URL http://ideas.repec.org/p/cep/cepdps/dp0749.html

Gaggl P, Wright GC. 2014. A short run view of what computers do: Evidence from a uk tax incentive. Technical report, UNC Charlotte.
URL http://belkcollegeofbusiness.uncc.edu/pgaggl/research/docs/UK_ICT_6-12-2014_WP.pdf

Garicano L, Rossi-Hansberg E. 2006. Organization and Inequality in a Knowledge Economy. *The Quarterly Journal of Economics* **121**: 1383–1435.
URL http://ideas.repec.org/a/tpr/qjecon/v121y2006i4p1383-1435.html

Garicano L, Rossi-Hansberg E. 2015. Knowledge-Based Hierarchies: Using Organizations to Understand the Economy. *Annual Review of Economics* **7**: 1–30.
URL https://ideas.repec.org/a/anr/reveco/v7y2015p1-30.html

Goldin C, Katz L. 2008. *THE RACE BETWEEN EDUCATION AND TECHNOLOGY*. Harvard University Press. ISBN 9780674028678.
URL http://books.google.com/books?id=mcYsvvNEUYwC

Goolsbee A. 1998. Investment Tax Incentives, Prices, And The Supply Of Capital Goods. *The Quarterly Journal of Economics* **113**: 121–148.
URL http://ideas.repec.org/a/tpr/qjecon/v113y1998i1p121-148.html

Hall RE, Jorgenson DW. 1967. Tax policy and investment behavior. *The American Economic Review* **57**: pp. 391–414. ISSN 00028282.
URL http://www.jstor.org/stable/1812110

House CL, Shapiro MD. 2008. Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation. *American Economic Review* **98**: 737–68.
URL http://ideas.repec.org/a/aea/aecrev/v98y2008i3p737-68.html

Imbens GW, Lemieux T. 2008. Regression discontinuity designs: A guide to practice. *Journal of Econometrics* **142**: 615–635.
URL https://ideas.repec.org/a/eee/econom/v142y2008i2p615-635.html

Ivus O, Boland M. 2015. The employment and wage impact of broadband deployment in Canada. *Canadian Journal of Economics* **48**: 1803–1830.
URL https://ideas.repec.org/a/cje/issued/v48y2015i5p1803-1830.html

Jaimovich N, Siu HE. 2012. The trend is the cycle: Job polarization and jobless recoveries. Working Paper 18334, National Bureau of Economic Research.
URL http://www.nber.org/papers/w18334

Jorgenson DW, Yun KY. 1990. Tax reform and u.s. economic growth. *Journal of Political Economy* **98**: pp. S151–S193. ISSN 00223808.
URL http://www.jstor.org/stable/2937635

Krueger AB. 1993. How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984-1989. *The Quarterly Journal of Economics* **108**: 33–60.
URL http://ideas.repec.org/a/tpr/qjecon/v108y1993i1p33-60.html

Lee DS, Card D. 2008. Regression discontinuity inference with specification error. *Journal of Econometrics* **142**: 655–674.
URL http://ideas.repec.org/a/eee/econom/v142y2008i2p655-674.html

Lee DS, Lemieux T. 2010. Regression Discontinuity Designs in Economics. *Journal of Economic Literature* **48**: 281–355.
URL https://ideas.repec.org/a/aea/jeclit/v48y2010i2p281-355.html

Lucas RJ. 1976. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy* **1**: 19–46.
URL http://ideas.repec.org/a/eee/crcspp/v1y1976ip19-46.html

Macaulay C. 2003. Job mobility and job tenure in the uk. *Labour Market Trends* **111**: 541–550.

Michaels G, Natraj A, van Reenen J. 2014. Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years. *The Review of Economics and Statistics* **96**: 60–77.
URL http://ideas.repec.org/a/tpr/restat/v96y2014i1p60-77.html