Industrial & Labor Relations Review

Volume 61, *Issue* 4 2008 *Article* 4

Returns to Pencil Use Revisited

Alexandra Spitz-Oener*

*Humboldt University, Berlin,

Returns to Pencil Use Revisited

Alexandra Spitz-Oener

Abstract

Many researchers believe that the observed positive association between computer use and wages simply reflects unobserved heterogeneity: like pencils and other "whitecollar" tools, computers are assigned to employees who possess productive attributes that would attract higher wages in any event. This study evaluates that claim by identifying the mechanisms through which computers changed the wage structure in West Germany in the late 1990s. The author finds that the spread of computers—but not of pencils—shifted the task composition of occupations toward analytical and interactive tasks that are complementary to computers' capabilities, and away from routine cognitive and manual tasks for which computers tend to substitute. Employees possessing computer-complementary skills enjoyed wage increases because computers both raised the demand for their skills and increased their marginal product.

KEYWORDS: computer use, wages

THE RETURNS TO PENCIL USE REVISITED

ALEXANDRA SPITZ-OENER*

Many researchers believe that the observed positive association between computer use and wages simply reflects unobserved heterogeneity: like pencils and other "white-collar" tools, computers are assigned to employees who possess productive attributes that would attract higher wages in any event. This study evaluates that claim by identifying the mechanisms through which computers changed the wage structure in West Germany in the late 1990s. The author finds that the spread of computers—but not of pencils—shifted the task composition of occupations toward analytical and interactive tasks that are complementary to computers' capabilities, and away from routine cognitive and manual tasks for which computers tend to substitute. Employees possessing computer-complementary skills enjoyed wage increases because computers both raised the demand for their skills and increased their marginal product.

The widespread use of computers is one of the fundamental changes in industrialized countries in recent decades. Today, there is a consensus that the implementation of computer technologies has changed skill requirements and hence the demand for labor. Because evidence suggests this change is skill-biased, computer use is also considered an important explanatory factor for the recent changes in the wage structure (see surveys by Katz and Autor 1999; Chennells and van Reenen 2002; Acemoglu 2002).

However, there are three competing explanations for the positive associations we observe between computer use and wages. First, employees who use computers at work may earn more because they are being rewarded for their computer skills. Second, it could be that even in the absence of computers, employees who use computers at work would earn higher wages than employees who do not. Computers are non-randomly assigned to employees, and hence the positive association might be due to unobserved heterogeneity. Third, individuals might be rewarded not specifically for using computers, but rather for performing tasks that are complementary to computers' functions. The increased diffusion of computers (owing to the exogenous decline in computer prices) then affects wages by increasing the demand for employees who possess the skills needed to perform the tasks to which computers are complementary. In addition, computers increase the marginal product of workers who use them to perform the complementary tasks, and therefore the computer users' wages increase.

*The author is Junior Professor of Economics, Humboldt University, Berlin. She thanks David Autor, Joshua Angrist, Irene Bertschek, Sandra Black, Francine Blau, Michael Burda, Bernd Fitzenberger, Alan Krueger, and Susanne Prantl for helpful comments and suggestions.

The data used in this article were obtained from the German Zentralarchiv fuer Empirische Sozialforschung at the University of Cologne (ZA). The data were collected by the Bundesinstitut fuer Berufsbildung and the Institut fuer Arbeitsmarkt-und Berufsforschung and documented by the ZA.

Copies of computer programs used to generate the tables in this paper are available from the author at alexandra.spitz-oener@wiwi.hu-berlin.de.

Industrial and Labor Relations Review, Vol. 61, No. 4 (July 2008). © by Cornell University.
0019-7939/00/6104 $01.00

This paper evaluates the third of these explanations by using a task-based approach, a framework that researchers now employ to investigate the detailed mechanisms by which computerization affects work (Autor, Levy, and Murnane 2003; Spitz-Oener 2006). Key to this framework is its conceptualization of work as a series of tasks, each of which can be characterized based on its substitutability and complementarity with computers. Because this framework provides a theoretical basis for identifying skills/tasks that are complementary to computers and an explanation for how the increased availability of computers increases the marginal product of employees performing these tasks, it allows me to reconcile the literature on wage-structure changes with that on premia for on-the-job computer use.

This paper uses the "Qualification and Career Survey," an individual-level data set from West Germany for 1979–99, to provide new evidence on the underlying causes of the return to computer use. Three features of the data set make it unique. First, it includes information on the activities that employees perform on the job, each of which falls into one of five task categories: non-routine analytic (such as researching and analyzing), non-routine interactive (such as managing and organizing), routine cognitive (such as calculating and bookkeeping), routine manual (such as operating machinery), and non-routine manual (such as serving and repairing). Second, it allows me to demonstrate that computer functions are complements to certain tasks (analytic and interactive non-routine tasks) and substitutes for others (manual and cognitive routine tasks). Finally, as the data set also includes information about office tools other than computers, it allows me to investigate what makes computers different from other tools. I focus in particular on pencil use, as pencils are the non-computer tool most fully discussed in the previous literature (see DiNardo and Pischke 1997).

Related Literature and Empirical Framework

Non-random assignment of computers to workers is the basis for the central critique of the literature on wage premia for on-the-job computer use. Previous empirical studies have addressed the issue in one form or another. All of the studies consider observed differences between computer users and non-users, although the number of observable variables has greatly differed across studies. The methods used to account for unobserved heterogeneity range from including proxies for individual ability in the regression specification to applying panel data methods.[1] The results are mixed, however, and depend largely on the underlying assumptions. The panel methods, for example, hinge crucially on the assumption of time-invariant unobserved heterogeneity. In the presence of changing returns to unobserved skills, differencing the data will not remove the wage effect of unobservables that might be correlated with computer use. As DiNardo and Pischke (1997) emphasized, this factor may be a plausible explanation for the differing results in panel analyses for the United Kingdom and France—because the wage structure has widened in the United Kingdom since the early 1980s but not in France.[2] Another argument that challenges the general credibility of panel estimates in this context is that identification in panel methods is through status changers, that is, individuals who started or stopped using a computer. In the presence of downwardly rigid wages, panel methods lead to an un-

[1]See Krueger (1993), who was the first to address the question of whether workers who use computers at work are paid more as a result of their computer skills. Studies using panel data for other countries are Entorf and Kramarz (1997) (France), Entorf, Gollac, and Kramarz (1999) (France), and Bell (1996) (the United Kingdom).

[2]Wage trends in Germany are often regarded as being similar to those in France (see, for example, Prasad 2004). However, Fitzenberger (1999), Fitzenberger, Hujer, McCurdy, and Schnabel (2001), and Dustmann, Ludsteck, and Schoenberg (2007) provide evidence that there have been changes in the wage structure in West Germany since the 1980s. The pattern of changes is actually quite similar to that observed in the United States, although the timing is different. The most commonly cited reason for the observed pattern is skill-biased technological change; changes in labor market institutions have played second-fiddle in most proposed explanations.

derestimation of the computer wage effect. Consider, for example, an employee who starts using a computer at work in one period and whose wage therefore increases. In the next period he stops using the computer but, as a result of the downward rigidity of wages, his wages do not decline. The panel method, which compares the wages of those who started using a computer with the wages of those who stopped using a computer, will then underestimate the wage effect of computer use.

Probably the most cited criticism of the "returns from computer use" literature comes from DiNardo and Pischke (1997), who showed that there is also a considerable wage effect from the use of pencils (and other "white-collar" tools) in cross-section estimates. If we do not believe pencils changed the wage structure, they argued, why should we believe computers did?

The task-based framework introduced by Autor et al. (2003) brings a new perspective to the literature on premia for on-the-job computer use. In this model, work is conceptualized as a series of tasks, each of which is classified as either routine or non-routine. Both manual and cognitive routine tasks are well-defined in the sense that they are easily programmable and can be performed by computers at economically feasible costs—a feature that makes routine tasks amenable to substitution by computer capital (Levy and Murnane 1996). Non-routine tasks, in contrast, are not well defined and programmable and, as things currently stand, cannot be easily accomplished by computers. However, computer capital is complementary to both analytical and interactive non-routine cognitive tasks in the sense that computer technology increases the productivity of employees performing these tasks.

This framework implies that there are two channels through which computerization affects wages. First, it shifts the content of work toward non-routine cognitive tasks (for which computers are complements) and away from manual and cognitive routine tasks (for which computers are substitutes), and therefore increases the demand for employees who possess the skills needed to perform non-routine cognitive tasks. (See Autor et al. [2003] and Spitz-Oener [2006] for detailed evidence on the shifts in the task composition of work.) Wage changes then depend on whether the supply of employees possessing these skills keeps up with the demand changes. The second channel through which computerization affects wages is more direct: the increased availability of computers increases the marginal product of employees who use computers to perform complementary tasks. The changes through both channels are triggered by the increased diffusion of computers at the workplace owing to the declining price of computer capital.[3]

In a competitive framework, one would expect the wage premium for computer use to be competed away. Efficiency wage considerations provide arguments for why the wage differential might be persistent over time (similar to inter-industry wage differentials or establishment wage differentials among workers with identical characteristics). Employers might, for example, find it profitable to pay computer users higher wages since such wage premiums might elicit effort when employers have only imperfect information concerning the behavior of workers on the job. Bresnahan, Brynjolfsson, and Hitt (2002), for example, discussed how the implementation of information technology in firms is associated with greater reliance on lateral communication, decentralized decision-making, and greater worker discretion.

The previous literature has mainly focused on the question of whether computer skills, that is, the skills enabling a worker to use a computer, are valuable in the labor market. Overall, the "returns to computer skills" approach is quite different from the task-based approach that focuses on how computerization affects the task/skill content of work, and whether employees who use a computer to perform certain tasks become more productive as a result. In this framework, the ability to use a computer is not valuable *per se;* rather, its value depends on the value of the tasks it enables a worker

[3]Autor et al. (2003) presented the general equilibrium model that is the foundation of this informal reasoning.

to perform. Borghans and ter Weel (2004) adopt a somewhat similar perspective when they contrast computer skills with other skills employees use at work, such as math or writing skills. However, the task-based framework suggests a different interpretation of their results of positive returns for the use of math and writing skills. Math, for example, is a routine task from a computer's perspective (in general, computers perform calculations much faster than humans can). However, deciding what kind of math one needs to solve a problem is an analytical task that, as yet, must be performed by a person (so the dummy for "math" in the wage regression is very likely to be a proxy for the analytical tasks employees perform at work).

Data

The analysis in this paper is based on the "Qualification and Career Survey," which is an employee survey carried out by the German Federal Institute for Vocational Training (Bundesinstitut für Berufsbildung, BIBB) and the Research Institute of the Federal Employment Service (Institut für Arbeitsmarkt-und Berufsforschung, IAB). For most of the analysis, I use the most recent cross-section, launched in 1998/99. It covers more than 30,000 men and women.[4] DiNardo and Pischke (1997) used 1979, 1985/86, and 1991/92 cross-sections from the same data set.

The survey contains information on monthly earnings, rank-ordered into 18 brackets. To calculate hourly wages for a worker, I divided the midpoint of the monthly earnings bracket of that worker by the worker's usual hours of work per month.[5] Unlike other data sets often used in wage analyses, such as the Current Population Survey for the United States, this data set has the advantage that earnings of highly paid workers are not censored from above. In all estimates, the logarithm of wages is used as the dependent variable. On average, employees in West Germany earned about 27 German Marks (about $16 U.S.) in 1998/99. Summary statistics are in Table 1.

An important, and unique, feature of this data set is its inclusion of information on the task composition of occupations; survey participants indicated what kinds of activities they performed at the workplace. Based on these activities, I construct five task categories: non-routine analytic tasks, non-routine interactive tasks, routine cognitive tasks, routine manual tasks, and non-routine manual tasks. Table 2 shows the list of activities that employees were asked about and how the activities are classified into the five task categories.[6] At the individual-level i, the task measures ($Task_{ik}$) are defined as

$$(1) \quad Task_{ik} = \frac{\text{number of activities in category } k \text{ performed by } i \text{ in 1998/99}}{\text{total number of activities in category } k \text{ in 1998/99}} * 100,$$

where $k = 1$: non-routine analytic tasks; $k = 2$: non-routine interactive tasks; $k = 3$: routine cognitive tasks; $k = 4$: routine manual tasks; and $k = 5$: non-routine manual tasks. For example, if the analytical task category includes 4 activities and employee i performs 2 of them, the analytical task measure for employee i is 50.

The data set also contains information about the employee's current occupation.

[4] I restrict the sample to West German residents with German nationality; in other words, East German residents and non-German employees are excluded from the sample. Also excluded are the self-employed, employees with agricultural occupations, employees working in the agricultural sector, and persons either under 18 or over 65 years of age.

[5] Comparable procedures have often been used in the literature. See, for example, DiNardo and Pischke (1997) and Entorf and Kramarz (1997).

[6] Instead of grouping the activities, one could include them individually in the wage regressions. Unreported results show, however, no statistically significant differences between the coefficients on the individual activities within task groups, except for "selling" (which is negatively related to wages) and "negotiating" (positively related to wages, with a larger coefficient than the other activities in the interactive task category). The size of other coefficients changed only marginally in response to the inclusion of the disaggregated tasks, and the level of statistical significance was not affected.

Table 1. Summary Statistics: West German Workers, 1998/99.

Independent Variable	*Mean*	*Std. Deviation*	*Min.*	*Max.*
Computer	0.57	0.50	0	1
Pencil	0.92	0.27	0	1
(Hourly) Wages (in German Marks)	27.19	11.82	3.13	98.68
Qualification				
High Education Level	0.17	0.37	0	1
Medium Education Level	0.71	0.46	0	1
Low Education Level	0.12	0.33	0	1
Experience	20.76	11.58	0	47
Tenure	11.75	9.84	0	47
Workplace Characteristics				
Non-Routine Analytic Tasks	14.01	23.80	0	100
Non-Routine Interactive Tasks	30.26	28.18	0	100
Routine Cognitive Tasks	21.73	41.24	0	100
Routine Manual Tasks	17.43	30.83	0	100
Non-Routine Manual Tasks	24.32	24.99	0	50
Company Characteristics				
Product Innovation	0.37	0.48	0	1
Process Innovation	0.51	0.50	0	1
Very Good Company Performance	0.18	0.39	0	1
Good Company Performance	0.65	0.48	0	1
Rather Bad Company Performance	0.14	0.35	0	1
Bad Company Performance	0.03	0.17	0	1
Other Controls				
Ever Unemployed	0.30	0.46	0	1
Married	0.69	0.46	0	1
Civil Servants	0.11	0.31	0	1
Born in East Germany	0.04	0.19	0	1
Woman	0.44	0.50	0	1
Lives in City	0.38	0.48	0	1

Source: Qualification and Career Survey.

Occupations are grouped according to the (2-digit) classification of occupational titles by the Federal Employment Bureau in 1999, leading to 78 occupational groups.

Another important feature of the data set is that it includes detailed information on the tools and machines used by the employees in the workplace. The "computer use" variable is a dummy that takes the value 1 if the employee used a computer, terminal, or electronic data-processing device on the job. The "pencil use" variable takes the value 1 if survey participants indicated that they used a writing tool at work.

I distinguish three levels of formal education attained by employees. Employees with a low level of education are those who had no vocational training. Employees with medium levels of education had a vocational qualification, either through an apprenticeship or through graduation from a vocational college. Employees holding a degree from a university or a technical college are classified as having a high level of education. As shown in Table 1, the majority of the survey participants, 70%, had a medium qualification level, whereas 17% were highly qualified and 12% had a low education level.

The survey participants also indicated their first year of work and the year they started to work for the current employer. Based on these answers, I calculate the years of (potential) work experience (1999 minus the first year of work) and tenure (1999 minus the first year with the current employer). The data set includes information about previous unemployment spells (a dummy variable: "Have you ever been unemployed before?"), marital

Table 2. Assignment of Activities.

Classification	*Tasks*
Non-Routine Analytic	researching, evaluating and planning, making plans, constructing, designing, sketching; working out rules/regulations; using and interpreting rules
Non-Routine Interactive	negotiating, lobbying, coordinating, organizing; teaching or training; selling, buying, advising customers, advertising; entertaining or presenting; employing or managing personnel
Routine Cognitive	calculating, bookkeeping; correcting of texts/data; measuring of length/weight/temperature
Routine Manual	operating or controlling machines; setting up machines
Non-Routine Manual	repairing or renovation houses/apartments/machines/vehicles; restoring art/monuments; serving or accommodating

status, gender, and whether survey participants were born in East Germany or were working as civil servants. I also constructed a dummy variable indicating residence in a city (a metropolitan area with more than 100,000 inhabitants).

One drawback of traditional individual-level data is that they generally do not provide information on employers. In contrast, the data set used in this study allows me to take various company characteristics into account, such as company size, industry affiliation, innovation strategy, and company performance. The ability to control for such variables represents a substantial improvement over many previous studies, given empirical evidence that, for example, larger companies and innovative companies pay higher wages and are more likely to use computers intensively (see, for example, Brown and Medoff 1989; Schmidt and Zimmermann 1991). The data set also includes information about company performance. The survey participants were asked whether the company was doing very well, well, rather badly, or badly. For each of these categories, I constructed a dummy variable. Table 1 shows that 18% of employees reported to work in companies that were doing very well and 65% worked in companies that were doing well. Seventeen percent of employees worked in companies that were doing either badly or rather badly.

Companies are classified according to 48 detailed industry codes. Based on these codes I group companies into three sectors: manufacturing, trade, and services.[7] The inclusion of these variables accounts for inter-industry wage differentials that are not already captured by the observed individual and company characteristics.

Empirical Results

The most cited criticisms of the "returns from computer use" literature come from DiNardo and Pischke (1997), whose cross-section estimates showed that computer use is not alone in being associated with considerable wage effects: the same association is found for pencils and other "white-collar" tools. Columns (1)–(3) of Table 3 recapitulate parts of DiNardo and Pischke's Table III. Panel A reports the coefficients for computer use when separate regressions are performed for each workplace tool. Panel B shows the results when all tools are included together in the specification. Controlling for the different workplace tools attenuates the coefficient for computer use in each period.

I reproduce DiNardo and Pischke's estimates using the most recent wave of the data set (Table 3, last column). Relative to the result in Panel A, the estimated coefficient for computer use decreases by about 30% due to the inclusion of the dummies for the different workplace tools (Panel B). The

[7]I also ran regressions that included more detailed industry dummies. The results reported in the next section are robust to this change in specification.

Table 3. OLS Regressions for the Effect of Different Tools on Pay.
Dependent Variable: Log(Hourly Wages)

Tool	*Germany, 1979*	*Germany, 1985–86*	*Germany, 1991–92*	*Germany, 1998–99*
	A. Tools Entered Separately			
Computer	.112	.157	.171	.204
	(.010)	(.007)	(.006)	(.006)
	B. Tools Entered Together			
Computer	.066	.105	.126	.146
	(.010)	(.008)	(.007)	(.007)
Calculator	.017	.053	.044	.051
	(.008)	(.007)	(.007)	(.006)
Telephone	.072	.043	.045	–.019
	(.007)	(.008)	(.008)	(.006)
Pen/Pencil	.062	.031	.035	.003
	(.007)	(.008)	(.008)	(.011)
Work while Sitting	.058	.050	—	.065
	(.007)	(.007)	—	(.006)

Standard errors are in parentheses. Source for columns (1)–(3): DiNardo and Pischke (1997), p. 299, Table 3, 2nd part, left side. Column (4): own regressions. Similar to the specification in DiNardo and Pischke, these regressions include education, experience, experience squared, and dummies for part-time, city, female, married, female*married, and civil servant.

coefficient for computer use still has a magnitude of about 15%. In contrast to previous years, in which the different tools have always had a statistically significant positive impact on wages, only the coefficients for the use of calculators and working while sitting remain positive and significant in 1998/99. The coefficient for using a pencil is statistically insignificant, and using a telephone at work now has a statistically significant negative correlation with wages. In what follows, I analyze this pattern in more detail.[8]

Columns (1)–(5) of Table 4 display the estimation results from "standard" wage regressions. Unreported results show that the raw log wage differential for computer use in Germany was 0.275 (about 31%) in 1998/99. This figure is slightly lower than the raw log wage differential of 0.288 that DiNardo and Pischke (1997) reported for Germany based on the 1991/92 cross-section of the BIBB/IAB data. Thus, in contrast to the period between 1979 and 1991/92, when the raw log wage differentials for computer use increased steadily (although at a declining rate), as shown in the paper by DiNardo and Pischke, this differential remained stable or even declined slightly in the 1990s.

Columns (1)–(5) show the results of specifications that are successively extended with additional controls. Including individual characteristics such as the level of formal education, work experience, tenure, gender, marital status, and residence in a city reduces the coefficient on computer use by more than 30% (column 1).[9] The coefficients on the controls have the expected signs. In column (2), workplace characteristics are included in the specification. The higher the measure for non-routine cognitive activities, both analytical and interactive, the higher the wages. In contrast, wages decrease in the measure for non-routine manual activities. Including the workplace characteristics additionally reduces the computer coefficient by 30% (compared to the coefficient in column

[8]The correlation between tools ranges from 0.007 (between work while sitting and pencil use) to 0.350 (between work while sitting and computer use).

[9]Note that in each specification of this table, the controls are also included interacted with a gender dummy.

Table 4. OLS Regressions for the Effect of Computers on Wages.
Dependent Variable: Log(Hourly Wages)

Variable	*(1)*	*(2)*	*(3)*	*(4)*	*(5)*	*(6)*	*(7)*
Computer Use	.185***	.126***	.123***	.123***	.073***	.131***	.078***
	(.005)	(.006)	(.008)	(.008)	(.010)	(.012)	(.027)
Non-Routine Analytic Tasks		.075***	.097***	.097***	.056***	.008	–.001
		(.014)	(.019)	(.019)	(.019)	(.040)	(.037)
Non-Routine Interactive Tasks		.153***	.151***	.151***	.132***	.176***	.124***
		(.013)	(.016)	(.016)	(.017)	(.029)	(.032)
Routine Cognitive Tasks		.024*	.012	.010	.009	.005	.010
		(.013)	(.015)	(.015)	(.016)	(.021)	(.022)
Routine Manual Tasks		–.006	–.048***	–.043**	–.037*	–.047*	–.036
		(.018)	(.020)	(.020)	(.022)	(.028)	(.032)
Non-Routine Manual Tasks		–.128***	–.105***	–.104***	–.089***	–.048*	–.041
		(.013)	(.017)	(.017)	(.017)	(.027)	(.027)
Non-Routine Analytic Tasks * Woman		.044*	.074**	.073**	.065**	.217***	.179**
		(.025)	(.034)	(.034)	(.035)	(.096)	(.089)
Non-Routine Interactive Tasks * Woman		–.011	.007	.008	.047	.041	.150***
		(.023)	(.028)	(.028)	(.030)	(.057)	(.061)
Routine Cognitive Tasks * Woman		.033	–.028	–.023	–.035	–.026	–.035
		(.029)	(.035)	(.035)	(.035)	(.049)	(.051)
Routine Manual Tasks * Woman		–.040	.024	.018	.001	.019	–.038
		(.039)	(.044)	(.044)	(.050)	(.063)	(.069)
Non-Routine Manual Tasks * Woman		.024	.026	.027	.015	.056	.016
		(.022)	(.029)	(.029)	(.031)	(.051)	(.055)
Individual Characteristics	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Company Characteristics	No	No	Yes	Yes	Yes	Yes	Yes
77 Occupation Dummies	No	No	No	No	Yes	No	Yes
7 Company Size Dummies	No	No	Yes	Yes	Yes	Yes	Yes
Dummies for Manufacturing & Trade	No	No	Yes	Yes	Yes	Yes	Yes
10 Dummies for West German States	No	No	No	Yes	Yes	Yes	Yes
Computer * $[x - E(x)]$	No	No	No	No	No	Yes	Yes
R^2	.353	.346	.401	.403	.452	.413	.469
Number of Observations	18,547	15,266	8,936	8,936	8,897	8,936	8,897

Employees with low levels of education working in large companies in the services sector are the base category. Controls interacted with a gender dummy are also included in the specifications. Heteroskedasticity-consistent standard errors are in parentheses in columns (1)–(5). Columns (6) and (7): *Computer* * $[x - E(x)]$ means that the specification includes interaction terms between computer use and all the level variables *X*, where all the *X* had previously been de-meaned by the sample averages. Bootstrapped standard errors using 1,000 resamples.

*Statistically significant at the 10% level; **at the 5% level; ***at the 1% level.

1). In column (3), company characteristics such as company size and information about the innovation strategy of the company are included in order to control, for example, for company size effects in remuneration. Industry dummies are also included to account for cross-sectoral differences in computer use and pay. In addition, dummies indicating company performance are included. The most interesting of the results from these exercises is the finding that although including company characteristics in the specification reduces the returns to education and eliminates the statistical significance of the dummies for birth in East Germany and city residence, it hardly affects the computer coefficient. Column (4) includes 10 dummies for West German states that control for cross-state differences in wage levels owing to, for example, differing economic conditions. These variables affect neither the computer coefficient nor the coefficients on the other controls. The specification in column (5) includes 77 two-digit occupation dummies. The occupation dummies have a large impact on the estimated computer wage differential.

A review of the results for the successive specifications shows that the computer coefficient drops by more than 70% between the raw logarithm wage differential and the specification represented by column (5); thus, the largest part of the raw logarithm wage differential for computer users has been due to observable differences that would have resulted in employees earning different wages even in the absence of computers. The results indicate that observable workplace characteristics such as workplace tasks account for the largest proportion of the bias in the raw logarithm wage differential. Conditional on all the controls, however, the results still suggest that employees who used computers on the job earned about 8% higher wages than employees who did not. Without occupation controls, the coefficient is even higher—0.123.[10]

Ideal control variables are those that are attributes solely of the assignment to computer use and the earnings process, and are unaffected by the treatment itself (for example, time-invariant individual characteristics such as gender and place of birth). Some of the controls used in this study, such as work experience and the incidence of previous periods of unemployment, might be affected by the treatment. Computer users are, for example, less likely than other workers to become unemployed. Therefore, the treatment effect estimated here does not capture the indirect effects of computer use on wages (for example, through productivity).

In columns (6) and (7) I use a less restrictive specification that includes interactions between all covariates X (previously de-meaned using sample averages) and the computer use dummy D_i. Hence, in contrast to the previous regressions, the coefficient on computer use is not constrained to be constant conditional on the observable variables. Although the previous specifications already include a large number of controls, the remaining 13 (8)% wage markup for computer users may still be due to characteristics that are not observable in the data set at hand if these unobservables are positively correlated with both computer use and wages. The aim of this approach, often termed regression-based matching or fully interacted linear matching, is to purge the specification of the remaining covariance between unobservables and computer use.[11] The specification in column (6) excludes occupation dummies, whereas the specification in column (7) includes them. The results in column (6) show that including the interaction terms increases the estimated coefficient of computer use to .131 (= average treatment effect, ATE). The interaction terms are jointly significant ($F_{(55,8824)} = 2.82$, $p = 0.0000$) and therefore provide evidence of the presence of heterogeneous effects.[12]

Table 5 shows the results for selected interaction terms. The empirical evidence suggests, for example, that computer users with a university degree or with high levels of analytical tasks benefited particularly in terms of wages. Gender, by contrast, did not have a significant effect on the gains from computer use. The estimated average treatment effect for the treated (ATT) is 0.142 (S.E. = 0.016) based on this regression, which is statistically significantly different from the ATE.[13] On av-

[10]Unreported results show that the main conclusions from these regressions do not change when a consistent sample of 8,936 observations is used (column 4).

[11]See Wooldridge (2002), Chap. 18.

[12]The standard errors in columns (6) and (7) are estimated using a design-matrix bootstrap approach (1,000 resamples). This method accounts for the fact that—before constructing the interaction terms—all the controls were de-meaned using the sample averages rather than the population averages. "Traditional" standard errors and test statistics are invalid because they ignore the sampling variation (this is referred to as the generated regressor problem; see Wooldridge 2002, Chap. 6).

[13]$\widehat{ATT} = \widehat{ATE} + (\sum_{i=1}^{N} D_i)^{-1}(\sum_{i=1}^{N} D_i)(X - \bar{X})\hat{\delta})$. $D_i = 1$ indicates that employee i uses a computer on the job, $D_i = 0$ that he or she does not. $D_i(X - \bar{X})$ are the interaction terms between computer use and all the level variables X, where all the X had previously been de-meaned by the sample averages. The $\hat{\delta}$ are the estimated coefficients of the interaction terms.

erage, treatment effect heterogeneity seems to be important. The result that the ATT is higher than the ATE suggests that there was selection into treatment based on expected returns. If, counterfactually, the group that started using computers had been the group that actually did not, the wage gain would have been substantially below that observed.[14]

Column (7) shows the results when occupation dummies are included as additional controls. This again reduces the estimated coefficient considerably, the ATE now being 0.078. The interaction terms are again jointly significant ($F_{(172,8535)} = 2.17$, $p = 0.0000$). The estimated ATT is 0.085 (S.E. = 0.036) based on this regression. As in column (6), the ATT and ATE are significantly different from each other.

As heterogeneous treatment effects appear to be important, the results might differ depending on the particular assumptions made, and assumptions on the out-of-sample predictions are particularly important in this respect. As a robustness test, I also did matching on the propensity score. Although regression-based matching and matching on the propensity score both address selection problems arising due to observable differences between computer users and non-users, they deal differently with another source of biases: differences in the distribution of observable characteristics for computer users and non-users (so-called common support problem).[15] Regression-based matching uses functional form assumptions to deal with the problem, while matching is typically performed only over the common support region.

Table 5. Selected Interaction Terms. Dependent Variable: Log(Hourly Wages)

High Educ. Level ∗ Computer	.124*
	(.075)
Medium Educ. Level ∗ Computer	–.053
	(.035)
Woman ∗ Computer	–.016
	(.086)
Non-Routine Analytic Tasks ∗ Computer	.115***
	(.047)
Non-Routine Interactive Tasks ∗ Computer	–.007
	(.035)
Routine Cognitive Tasks ∗ Computer	.023
	(.031)
Routine Manual Tasks ∗ Computer	–.000
	(.041)
Non-Routine Manual Tasks ∗ Computer	–.104***
	(.035)
Experience ∗ Computer	.007**
	(.003)
Experience^2 ∗ Computer	–.000*
	(.000)

Notes: These results are based on the specification in Table 4, column (6). None of the coefficients for the variables interacted with a gender dummy were statistically different from those for men.

*Statistically significant at the 10% level; **at the 5% level; ***at the 1% level.

Table 6 displays the results of the matching estimations when occupations are not included in the specification estimating the propensity score. Column (1) shows that, using nearest neighbor matching, the average wage effect of computer use for employees who use computers is 15%. The coefficient is significant at the 1% level. In the matching, 3,461 controls are used to estimate the wages computer users would have earned had they not adopted computer technology.[16] Owing to the restriction that only observations within the region of common support are used, 400 computer users are not in the analysis. Column (2) shows the results when an Epanechnikov kernel is used in the matching process, which slightly lowers the magnitude of the ATT.

[14]By focusing mainly on the ATE, previous studies have neglected the potential for parameter heterogeneity. In a recent contribution, Dolton and Makepeace (2004) provided evidence for the importance of parameter heterogeneity by allowing for variation in the parameter values for different types of computer users—those who used computers in both periods considered (stayers), only in the first period (leavers), and only in the second period (enterers). However, they restricted the parameters to be homogeneous across other characteristics. My approach of estimating the ATT and allowing for parameter heterogeneity across all characteristics is more general.

[15]See, for example, Heckman, Ichimura, Smith, and Todd (1998) for a discussion of the different bias components.

[16]Owing to missing values in the matching variables, the sample reduces to 9,240 individuals (5,779 computer users and 3,461 computer non-users) in the matching specification.

Table 6. Results of the Propensity Score Matching, with and without Occupation Dummies. Dependent Variable: Log(Hourly Wages)

	(1)	*(2)*	*(3)*	*(4)*
A. Without Occupation Dummies				
Computer	.137***	.133***	.146***	.138***
	(.028)	(.032)	(.036)	(.024)
B. With Occupation Dummies				
Computer	.078***	.077***	.062***	.082***
	(.038)	(.021)	(.025)	(.022)
Number of Treated	5,379	5,379	5,076	5,076
Number of Controls	3,461	3,461	3,461	3,461

Notes. Column explanations: (1) ATT of computer use using nearest neighbor matching (random draw version, bootstrapped standard errors using 1,000 replications). (2) ATT using kernel matching (Epanechnikov kernel, bootstrapped standard errors using 1,000 replications). (3) ATT of computer use using nearest neighbor matching (random draw version, bootstrapped standard errors using 1,000 replications) with the additional restriction that treated and controls have the same level of education. (4) ATT using kernel matching (Epanechnikov kernel, bootstrapped standard errors using 1,000 replications) with the additional restriction that treated and controls have the same level of education. Only observations that are on the common support are used. The caliper is set to .001. The propensity score is estimated using the level of formal education, age, age^2, work experience, work experience^2, interaction between work experience and education, workplace tasks, born in East Germany, ever unemployed, living in a city, woman, married, married woman, 8 company size dummies, 39 industry dummies, 3 dummies reflecting company performance, 10 dummies for West German states, and a constant as regressors; the Panel B estimation also includes 77 occupation dummies.

*Statistically significant at the 10% level; **at the 5% level; ***at the 1% level.

The last column of Table 7 shows the means of the main individual characteristics of the computer non-users that are used as controls in the matching model. It is evident that, although there is a convergence between treatment subjects and controls with respect to most of the observable characteristics, the difference in education level is still statistically significant. Therefore, columns (3) and (4) of Table 6, Panel A, show estimates from specifications in which matches had to exhibit not only similar propensity scores, but also equal levels of education. This extension increases the number of computer users who are off the common support to 703. The estimate from nearest neighbor matching increases to about 16%, whereas the coefficient on the estimate that uses an Epanechnikov kernel increases to 15%. Both estimates are highly statistically significant.

Panel B of Table 6 shows the results when occupation dummies are included in the specification that estimates the propensity score. Both Table 6's lower panel and Table 7 show that the estimated ATTs are very similar to those found when we use regression-based approaches (Table 4, columns 6 and 7, respectively). This may be the case because (i) there is no common support problem, (ii) there is little heterogeneity in treatment effects or all the propensity scores are small, and (iii) there is no serious mis-specification in the no-treatment outcome (see Blundell, Dearden, and Sianesi 2005).

Figure 1 shows the kernel density estimates of the distribution of propensity scores for computer users and non-users. The two distributions greatly overlap, and the propensity scores in the analysis assume values between 0 and 1. The mean is 0.64 and the median is 0.73. This information does not suggest that propensity scores are particularly clustered at certain points in the distribution. The results of the matching procedure reveal that 400 (703) of the 5,779 computer users are outside the region of common support. Overall, in this application, imputing the values that are outside the common support by relying on a functional form assumption or discarding the computer users without similar counterfactuals from the analysis does not lead to different results.

I now consider the DiNardo and Pischke (1997) criticism and estimate regressions that

Table 7. Mean Comparison for Computer Users and Computer Non-Users.

Independent Variable	*Computer User*	*Computer Non-User Prior to Matching*	*Selected Controls*†
High Education Level	.25	.07**	.17**
Medium Education Level	.69	.72**	.74**
Low Education Level	.06	.21 **	.09**
Age	39.99	4.29**	4.20
Experience	2.04	21.52**	21.22**
Tenure	12.37	1.82**	12.33
Ever Unemployed	.26	.34**	.29
Married	.70	.65**	.70
Woman	.44	.44	.44
Civil Servants	.15	.05**	.13
Born in East Germany	.03	.05**	.04

**Means differ with statistical significance of 5% in a two-tailed t-test between computer users and either computer non-users prior to matching (column 3) or the selected computer non-users (last column).

†Computer non-users who are selected by the matching procedure.

include dummies indicating the use of various workplace tools other than computers. The results for pencil use alone are shown in Table 8. Unreported results show that the first-order relationship between pencil use and wages is 5.7% (significant at the 1% level). Repeating the procedure used to generate the results in Table 4, I successively augment the specifications with additional controls. In contrast to the computer effect, the estimated pencil effect disappears as soon as controls for individual and workplace characteristics are included in the specification (column 2). In fact, the pencil use variable loses all statistical significance following inclusion of just a handful of the controls that, even *in toto,* had only attenuating effects on the computer use coefficient. Column (2) shows the results of the regression specification including the minimum number of controls that are necessary to eliminate the wage effect of pencil use. It is noteworthy that the controls for workplace tasks are what made the positive association between pencils and wages disappear; this result suggests that "pencils" picked up the positive association between analytical and interactive tasks and wages in the previous specification.

The question that now naturally arises is, which features of computers distinguish them from pencils? In this paper, the task-based framework is used to get a new perspective on this issue.

Autor et al. (2003) and Spitz-Oener (2006) contended that the "traditional" skill-biased technological change hypothesis is probably an insufficiently nuanced view of changes in the workplace associated with the recent diffusion of computer technologies. The argument is that computers substitute for certain tasks—those that are expressible in rules and thus programmable (termed routine tasks)—whereas they complement workers in performing nonroutine analytic and nonroutine interactive tasks. In contrast to this reasoning, I expect pencils neither to substitute for nor to complement certain tasks.[17]

For this part of the analysis I use all four waves of the data set, that is, 1979, 1985/86, 1991/92, and 1998/99. Table 9 shows how changes both in computer use (Panel A) and in pencil use (Panel B) are related to changes in the task composition of occupations in recent decades.[18] Each column

[17]DiNardo and Pischke (1997) noted that the computer coefficient increases strongly over time when all tools are entered in the regression together, while some of the coefficients on the other tools tend to fall over time. This result, they concluded, might indicate that the role of computers in the workplace is changing. While this view is plausible inasmuch as changes in computer technology itself (in particular, the convergence of information and communication technologies) have certainly altered its role in the workplace, the explanation I outline in this paper investigates changes in occupational skill requirements owing to computerization.

[18]Panel A is from Spitz-Oener (2006), Table 6.

Figure 1. Kernel Density Estimations of Propensity Scores for Computer Users and Computer Non-Users.

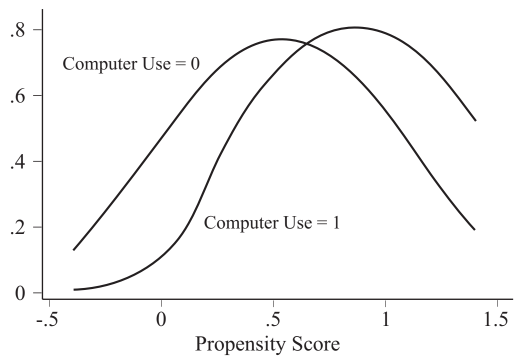

represents a separate OLS regression of the annualized changes in occupational task inputs on the annual changes in occupational computer use or the annual changes in occupational pencil use, respectively. Annual changes are calculated between successive waves, that is, between 1979 and 1985/86, between 1985/86 and 1991/92, and between 1991/92 and 1998/99. The regressions are then performed on the stacked data set. The specifications all include time dummies for 1985/86–1991/92 and 1991/92–1998/99 in order to control for within-occupation trends in task changes for the corresponding time period relative to the base period 1979–1985/86.

While the Panel A results indicate that growth in computerization was positively associated with growth in analytical and interactive task inputs and with declines in routine manual and routine cognitive task requirements, the results in Panel B tell a quite different story. Of the five equations, two—for the relationship between changes in *pencil* use and changes in (a) analytic and (b) routine manual activities—yield statistically insignificant coefficients. Statistically significant coefficients are found for interactive tasks (a negative relationship with changes in pencil use), nonroutine tasks (a negative relationship), and manual tasks (a positive relationship). Thus, it turns out that these occupations shifted away from interactive tasks toward more routine cognitive tasks.

An inspection of underlying occupations shows that computer users and pencil users belonged to a very similar set of occupations in 1979. Examples are office workers, technicians, and sales clerks. By 1998/99, however, these two groups were working in very different occupations. Office workers, technicians, and sales clerks remained among the typical occupations of computer users in 1998/99, as in 1979, but pencil users, in the later years, were more frequently crafts workers, such as mechanics, electricians,

Table 8. OLS Regressions for the Effect of Pencil Use on Wages. Dependent Variable: Log(Hourly Wages)

Description	*(1)*	*(2)*	*(3)*	*(4)*	*(5)*
Pencil Use	.033***	.016	.016	.017	.006
	(.010)	(.011)	(.013)	(.013)	(.013)
Individual Characteristics	Yes	Yes	Yes	Yes	Yes
Workplace Characteristics	No	Yes	Yes	Yes	Yes
Company Characteristics	No	No	Yes	Yes	Yes
77 Occupation Dummies	No	No	No	No	Yes
7 Company Size Dummies	No	No	Yes	Yes	Yes
Dummies for Manufacturing & Trade	No	No	Yes	Yes	Yes
10 Dummies for West German States	No	No	No	Yes	Yes
R^2	.304	.328	.385	.387	.447
Number of Observations	15,951	13,986	8,037	8,037	8,001

Notes: Employees with low levels of education working in large companies in the services sector are the base category. Controls interacted with a gender dummy are also included in the specifications. Heteroskedasticity-consistent standard errors are in parentheses.

*Statistically significant at the 10% level; **at the 5% level; ***at the 1% level.

Table 9. Changes in Computer Use, Pencil Use, and Skill Requirements.

	Dependent Variables: Annualized Changes in Task Inputs				
Independent Variable	*Analytic Tasks*	*Interactive Tasks*	*Routine Cognitive Tasks*	*Routine Manual Tasks*	*Non-Routine Manual Tasks*
			Panel A		
Δ Computer Use	.086***	.188***	–.312***	–.561***	.128
	(.032)	(.031)	(.105)	(.148)	(.085)
Dummy 1985/86–1991/92	–6.160***	3.536**	–1.960	–2.462	–1.67***
	(1.129)	(1.767)	(3.098)	(7.712)	(.401)
Dummy 1991/92–1998/99	–7.987***	8.915***	16.394**	–7.436	–1.147**
	(1.381)	(1.440)	(7.726)	(7.065)	(.520)
R^2	.183	.337	.079	.131	.065
Number of Observations			237		
			Panel B		
Δ Pencil Use	.025	–.112***	.913***	.074	–.320***
	(.025)	(.023)	(.130)	(.059)	(.077)
Dummy 1985/86–1991/92	–.604***	.677***	–1.987***	–.750	–1.036***
	(.139)	(.195)	(.407)	(.782)	(.424)
Dummy 1991/92–1998/99	–.821***	1.488***	–2.106***	–1.522**	.174
	(.202)	(.165)	(.829)	(.699)	(.591)
R^2	.150	.279	.431	.028	.128
Number of Observations			237		

Source, Panel A: Spitz-Oener (2006), Table 6. Robust standard errors are in parentheses; regressions are weighted by the number of individuals within the occupation group.

*Statistically significant at the 10% level; **at the 5% level; ***at the 1% level.

and masons, or health care workers, such as nurses and masseurs.

In combination with the results shown in Tables 4 and 5, these findings suggest that computers changed the wage structure by increasing the demand for analytical and interactive tasks that were positively associated with wages (Table 4, column 1–5). In addition, the numbers in columns (6) and (7) of Table 4, together with the results in Table 5, show that employees who used a computer to perform analytical tasks earned higher wages. Interestingly, the results for these unrestricted specifications indicate that in order to realize higher wages with a higher level of analytical task inputs, men—but not women—had to make use of computers at work.[19]

Overall, the findings are consistent with the predictions from the task-based framework that there are two channels through which computerization affects wages. First, it changes the content of work, with resultant upward pressure on the wages of employees who possess the skills to perform the tasks for which greater need has arisen and downward pressure on the wages of those whose previous job tasks are among those taken over by computers. The evolution of wages then depends on how the supply of these skills reacts to the change in demand. Second, and more directly, computers increase the marginal product of employees who perform analytical tasks, which explains the positive association between computer use, analytical tasks, and wages that we observe in the data.

Conclusions

Critics have dismissed much research on the returns to computer use on the grounds that the results are driven by unobserved heterogeneity. In the present study I have employed the task-based framework, an approach used to analyze the detailed mechanisms by which computers affect work, to provide a new perspective on this issue.

[19]Other work of mine examines the differences in task changes between men and women (Black and Spitz-Oener 2007).

I first show how computerization shifted the task content of work toward analytical and interactive tasks and away from cognitive and manual routine tasks. To address long-standing concerns raised by DiNardo and Pischke (1997), I then show that no such pattern is found for the spread of pencils on the job. These results, combined with the results of wage regressions of the kind that have typically been performed in the previous literature on the returns to computer use, suggest that a considerable fraction of the computer wage premium is actually attributable to the analytical and interactive tasks employees perform, variables that have not been observable in previous studies. As such, the positive association between computers and wages does not reflect a return to computer skills as human capital theory would typically predict, but a return to interactive and analytical skills that are associated with computer use. In addition, the results are consistent with the idea that computers increase the marginal product of employees performing certain tasks, which leads to higher wages.

While the results do not rule out alternative explanations, this paper suggests that previous discussions of computer wage premia have been misleading because they have applied an undifferentiated view of the impact of computers on the labor market. By providing a more nuanced view of the relationship between computers and wages, this paper takes the ongoing inquiry one important step further.

REFERENCES

Acemoglu, Daron. 2002. "Technical Change, Inequality, and the Labor Market." *Journal of Economic Literature,* Vol. 40, No. 1 (March), pp. 7–72.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics,* Vol. 118, No. 4 (November), pp. 1279–1333.

Bell, Brian D. 1996. "Skill-Biased Technical Change and Wages: Evidence from a Longitudinal Data Set." Technical Report, Nuffield College.

Black, Sandra E., and Alexandra Spitz-Oener. 2007. "Explaining Women's Success: Technological Change and the Skill Content of Women's Work." Working Paper 13116, National Bureau of Economic Research; forthcoming in *Review of Economics and Statistics.*

Blundell, Richard, Lorraine Dearden, and Barbara Sianesi. 2005. "Evaluating the Impact of Education on Earnings in the UK: Models, Methods, and Results from the National Child Development Survey." *Journal of the Royal Statistical Society,* Series A, Vol. 168, No. 3, pp. 473–512.

Borghans, Lex, and Bas ter Weel. 2004. "Are Computer Skills the New Basic Skills? The Returns to Computer, Writing, and Math Skills in Britain." *Labour Economics,* Vol. 11, No. 1 (February), pp. 85–98.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics,* Vol. 117, No. 1 (February), pp. 339–76.

Brown, Charles, and James Medoff. 1989. "The Employer-Size Wage Effect." *Journal of Political Economy,* Vol. 97, No. 5 (October), pp. 1027–59.

Chennells, Lucy, and John van Reenen. 2002. "The Effects of Technical Change on Skills, Wages, and Employment: A Survey of the Microeconometric Evidence." In Nathalie Greenan, Yannick L'Horty, and Jacques Mairesse, eds., *Productivity, Inequality, and the Digital Economy,* pp. 175–225. Cambridge, Mass.: MIT Press.

DiNardo, John E., and Joern-Steffen Pischke. 1997. "The Returns to Computer Use Revisited: Have Pencils Changed the Wage Structure Too?" *Quarterly Journal of Economics,* Vol. 112, No. 1 (February), pp. 291–303.

Dolton, Peter, and Gerry Makepeace. 2004. "Computer Use and Earnings in Britain." *Economic Journal,* Vol. 114, No. 494, pp. C117–C129.

Dustmann, Christian, Johannes Ludsteck, and Uta Schoenberg. 2007. "Revisiting the German Wage Structure." Discussion Paper 2685, IZA.

Entorf, Horst, Michel Gollac, and Francis Kramarz. 1999. "New Technologies, Wages, and Worker Selection." *Journal of Labor Economics,* Vol. 17, No. 3 (July), pp. 464–91.

Entorf, Horst, and Francis Kramarz. 1997. "Does Unmeasured Ability Explain the Higher Wages of New Technology Workers?" *European Economic Review,* Vol. 41, No. 8, pp. 1489–1509.

Fitzenberger, Bernd. 1999. "Wages and Employment across Skill Groups." *ZEW Economic Studies,* Vol. 6. Heidelberg and New York: Physica-Verlag.

Fitzenberger, Bernd, Reinhard Hujer, Thomas McCurdy, and Reinhold Schnabel. 2001. "Testing for Uniform Wage Trends in West-Germany: A Cohort Analysis Using Quantile Regressions for

Censored Data." *Empirical Economics,* Vol. 26, No. 1, pp. 41–86.

Heckman, James, Hidehiko Ichimura, Jeffrey Smith, and Petra E. Todd. 1998. "Characterizing Selection Bias Using Experimental Data." *Econometrica,* Vol. 66, No. 5 (September), pp. 1017–98.

Katz, Lawrence F., and David H. Autor. 1999. "Changes in the Wage Structure and Earnings Inequality." In Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics,* Vol. 3, pp. 1463-1555. Amsterdam: Elsevier Science.

Krueger, Alan B. 1993. "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989." *Quarterly Journal of Economics,* Vol. 108, No. 1 (February), pp. 33–60.

Levy, Frank, and Richard J. Murnane. 1996. "With What Skills Are Computers a Complement?" *American Economic Review,* Vol. 86, No. 2 (May), pp. 258–62.

Prasad, Eswar S. 2004. "The Unbearable Stability of the German Wage Structure: Evidence and Interpretation." *Staff Papers,* Vol. 51, No. 2, IMF.

Schmidt, Christoph M., and Klaus Zimmermann. 1991. "Work Characteristics, Firm Size, and Wages." *Review of Economics and Statistics,* Vol. 73, No. 4 (November), pp. 705–10.

Spitz-Oener, Alexandra. 2006. "Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure." *Journal of Labor Economics,* Vol. 24, No. 2 (April), pp. 235–70.

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data.* Cambridge, Mass.: MIT Press.