sociological science

The Rise of Programming and the Stalled Gender Revolution

Siwei Cheng, Bhumika Chauhan, Swati Chintala

New York University

Abstract: Despite remarkable progress toward gender equality over the past half-century, the stalled convergence in the gender wage gap after the mid-1990s remains a puzzle. This study provides new insights into this puzzle by conducting the first large-scale investigation of the uneven impact of the rise of programming in the labor market for men and women since the mid-1990s. We argue that the increasing reliance on programming has favored men's economic status relative to women's and therefore may help explain the slow convergence of the gender wage gap. We differentiate between two effects: (1) the composition effect, wherein men experience a greater employment growth in programming-intensive occupations relative to women, and (2) the price effect, wherein the wage returns to programming intensity increase more for men than women. Our empirical analysis documents a strong relationship between the rise of programming and the slow convergence of the gender wage gap among college graduates. Counterfactual simulations indicate that the absence of the composition and price effects would have reduced the gender wage gap over the past two decades by an additional 14.70 percent. These findings call attention to the role gender institutions play in shaping the uneven labor market impact of technological change.

Keywords: technological change; labor market; gender inequality; computerization

Citation: Cheng, Siwei, Bhumika Chauhan, and Swati Chintala. 2019. "The Rise of Programming and the Stalled Gender Revolution." Sociological Science 6: 321-351.

Received: December 16, 2018

Accepted: March 6, 2019

Published: April 30, 2019

Editor(s): Jesper Sørensen, Kim Weeden

DOI: 10.15195/v6.a13

THE narrowing of the gender wage gap is one of the most remarkable advances of the past half-century of the American labor market. However, since the mid-1990s, in contrast to the continued rise in women's college completion rates, the movement toward gender pay equity has slowed down and even stalled (Blau and Kahn 2006; England 2010; Gerson 2009; Goldin 2002). Various explanations have been proposed to address this puzzle, including persistent occupational and job segregation by gender (Charles and Grusky 2005; Tomaskovic-Devey and Skaggs 2002), the growing prevalence of overwork (Cha and Weeden 2014), deep-rooted gender essentialism and status beliefs (Correll 2004; Ridgeway and Correll 2004), continued discrimination in hiring and promotion (Hirsh 2009; Petersen and Saporta 2004), and persistent gender division of housework and childcare (Bianchi et al. 2012).

One important piece in this puzzle that has received relatively little attention is the uneven impact of technological change for men and women. The spread of computers in the workplace in the 1980s has been shown to affect relative wages in the workforce (Card and DiNardo 2002; Kristal 2013). Yet, the impact of computerization can differ substantially by gender. For instance, increased computer usage in the 1980s benefited the economic standing of women relative to men both because women are more likely to use computers at work and because computers restructured the workplace in a manner that deemphasized physical strength (Black and Spitz-Oener 2010; Blau and Kahn 2000; Krueger 1993; Weinberg 2000). In the meantime, the gender wage gap converged remarkably in the 1980s, but the convergence slowed down as the diffusion of computers reached a plateau in the early 1990s (Blau and Kahn 2006; Friedberg 2003).

The 1990s saw a twist in this story: Whereas the spread of general computer usage reached a plateau, new waves of advances in computer technology took off. Programming-related tasks and skills, such as coding and large-scale database processing, started creating new edges of wage inequality (Acemoglu and Restrepo 2017; Autor 2015). As the rise of programming took place against a backdrop of continued underrepresentation of women in science and technology occupations and their limited access to higher-ranked job positions in these occupations, this recent wave of technological changes may have impeded, rather than facilitated, the convergence of the gender wage gap since the mid-1990s. Although an emerging literature has started to examine the consequences of this recent technological change for overall inequality, no prior work has taken advantage of large-scale, long-term, nationally representative data to investigate its implications for the trend in the gender wage gap.[1]

This study offers new insights into the stalled gender revolution by investigating the impact of the post–mid-1990s rise in programming on the trend in the gender wage gap. Specifically, we differentiate between two effects. The first effect operates through a gender differential in changes in occupational composition; that is, men may have experienced a larger growth of employment in programming-intensive occupations than women, resulting in a growing underrepresentation of women in programming-intensive occupations. This compositional change, when coupled with the positive net wage returns associated with programming intensity, may impede the convergence of the gender wage gap. The second effect works through a gender differential in the trend of price, or the wage payoff, associated with programming intensity. That is, even without gender differences in occupational composition, a greater increase in the wage payoff to programming for men than for women, in itself, can serve to slow down the closing of the gender wage gap.

Our empirical investigation consists of four steps. Step 1 constructs measures of programming intensity and general computer usage intensity for detailed occupations. In step 2, we merge these occupation-level measures with individual-level data from 1994 to 2015 to analyze the trends in employment composition with regard to computer-related tasks and skills by gender and education. This step allows us to examine the composition effect. In step 3, we estimate the gender-specific net wage returns to programming intensity across two decades, controlling for a large set of demographic, educational, employment, and occupational characteristics. This step allows us to examine the price effect. Finally, in step 4, we use counterfactual simulations to evaluate the extent to which the composition and price effects impeded the convergence of the gender wage gap.

Our results demonstrate a strong relationship between the rise of programming and the slow convergence of the gender wage gap among college graduates. From the mid-1990s onward, among college graduates, there has been a continued increase in programming intensity in men's, but not women's, occupations. The wage premium associated with higher programming intensity has increased steadily for men over time while remaining unchanged for women. Counterfactual decompositions suggest that among college graduates, the composition effect has impeded the convergence of the gender wage gap by about 11.76 percent, and the price effect by about 14.18 percent. These findings suggest that the labor market impact of the rise of programming is highly gendered.

Theoretical Perspectives on the Labor Market Impact of Technological Change

There are two major perspectives on the impact of technological change on labor market inequality. The first focuses on the market demand for skills and tasks. The skill-biased technological change (SBTC) literature emphasizes that changes in technology affect the relative demand for, as well as the wage payoff to, various skills (Autor, Katz, and Krueger 1998; Goldin and Katz 1998; Katz and Murphy 1992; Levy and Murnane 1992). The labor market impact of computerization is one of the most prominent examples of the SBTC theory (Freeman 2002; Friedberg 2003; Krueger 1993). Krueger (1993) found that workers who use computers at work earn a 10- to 15-percent wage premium. A related line of work emphasizes the impact of technological change on the allocation of tasks (Acemoglu and Autor 2011; Autor 2013). Whereas skills are a worker's endowment of capabilities to perform various tasks, the task itself is the unit of concrete work activity that produces the output. As an increasing share of tasks are being standardized and automated, routine nonmanual tasks are reallocated from human labor to machines. At the same time, the demand for certain nonroutine tasks, particularly those involving coding and programming machines, is likely to grow (Brynjolfsson and McAfee 2014).

The market-based perspective above largely assumes that the impact of technological change can be boiled down to the market demand for skills and tasks. It may not, however, fully explain how the rise of programming affects men's and women's employment and wage outcomes differently. In contrast, a second perspective contends that technological changes may interact with nonmarket forces, such as social, institutional, and organizational factors, to influence inequality along multiple dimensions (Burris 1998; Charles and Grusky 2005; Kristal 2013; Kristal and Cohen 2015; Salzman and Rosenthal 1994; Weeden 2002). Social institutions can foster gatekeeping practices among already advantaged workers and thereby perpetuate existing between-group inequality. They can also condition how technological change affects micro-level processes surrounding individuals and the organizational context of their workplace. In the next section, we take the social and institutional perspective on technological change as the point of departure and discuss how the recent rise of programming may unevenly affect men and women in the labor market.

Linking the Rise of Programming and Trends in Gender Inequality

To understand the gendered impact of the rise of programming, we differentiate between two types of effects. Composition effect describes a situation wherein men experience a greater employment growth in programming-intensive occupations relative to women. Price effect operates when the wage returns to programming intensity increase more for men than for women. We show that both effects have impeded the convergence of the gender wage gap. In what follows, we draw on various theoretical perspectives to discuss these two effects. Here, we note that our primary goal is to document the impact of the rise of programming on the trend in gender inequality rather than to test specific mechanisms. Hence, the following discussion mainly serves to provide the theoretical basis for our empirical analysis.

Gender Differences in Occupational Composition Trends

The rise of programming in the labor market may impede the movement toward gender equality by affecting the relative representation of men and women in occupations with higher and lower levels of programming intensity. The literature indicates several intervening mechanisms. First, the theory of labor queues views the labor market as one in which occupations are ordered in terms of their attractiveness, and potential employees are ordered in a queue according to employers' ranking (Reskin and Roos 2009; Thurow 1975). When the demand for certain occupations grows, the newly opened opportunities will be allocated first to individuals or social groups placed at the front of the labor queue. Whereas the human capital approach views labor queues as determined primarily by potential workers' productivity (Becker 1957), the sociological perspective emphasizes social and cultural factors that affect employers' perceptions about who is more suitable for the job (Reskin and Roos 2009). Essentialist beliefs about women's intrinsic interests, abilities, and relative status in the workplace, for example, can lead to gender bias and discrimination in hiring (Charles and Grusky 2005; England 2010; Levanon and Grusky 2016; Ridgeway and Correll 2004). Applied to the case of programming-intensive jobs, this perspective suggests that employers may hold men as intrinsically better suited than women for jobs involving interaction with complex computer systems, writing code, and deriving mathematical formulations. This leads to the expectation that, when the demand for programming-related occupations grew, men were ranked at the front of the labor queue for taking these jobs, whereas the progress of women into these occupations remained largely stagnant.

Second, social closure mechanisms may also contribute to the diverging trends in employment in programming-intensive occupations between men and women. This perspective emphasizes the institutional barriers established around occupations that limit individuals or social groups from entering these occupations (Abbott 1988; Weeden 2002). As occupations that rely heavily on programming tend to be male dominated, current practitioners in these occupations—mostly men—may create social and institutional barriers that restrict women's access to these jobs, contributing to women's decreasing representation in these occupations.

Third, the gender difference in employment can also be driven by supply-side factors. The "leaks" in the pipeline carrying women to programming-intensive jobs may start early on, when boys are more likely than girls to attend computer camps (Hess and Miura 1985; Lapan et al. 2000; Wilder, Mackie, and Cooper 1985), and continue during high school and college, when biased expectations and evaluations in the cultural environment discourage female students in science, technology, engineering, and mathematics (STEM) fields in general (Correll 2001; Legewie and DiPrete 2014; Moss-Racusin et al. 2012) and computer science in particular (Cheryan et al. 2009; Cohoon 2001; Frenkel 1990). Whereas the gender gap in high school students' participation in math and science courses has narrowed over time, the proportional representation of women in STEM majors in college has stalled or even declined (England and Li 2006; Mann and DiPrete 2016; Xie, Fang, and Shauman 2015). Because programming-related jobs tend to require training in STEM fields, these gender differences on the supply side suggest that men are on average more likely than women to take advantage of newly available programming-related employment opportunities over time.

Moreover, the social and cultural processes shaping employers' preferences may also shape the psychological processes through which individuals come to form perceptions and expectations about themselves (Cech et al. 2011; Charles and Bradley 2009; Cooper and Weaver 2003; Correll 2001; Ridgeway 1997). For instance, recent evidence suggests that women often lack a sense of belonging to the culture of technology jobs, where masculinity and heteronormativity are the norm (Alfrey and Twine 2017; Rosenbloom et al. 2008; Wynn and Correll 2017). Women in high-technology firms also exhibit significantly higher attrition rates than men (Hill, Corbett, and St. Rose 2010).

Lastly, women's progress into programming-intensive occupations may be affected by the gendered division of family responsibilities, which limits women's access to jobs that require commitment to longer and less flexible work hours (Cha 2010; Hochschild and Machung [1989] 2012; Jacobs and Gerson 2001). This effect may be particularly strong in programming-intensive occupations, which tend to require commitment to long work hours individually and in teams (Freeman 2002; Perlow 1999). Additionally, these domestic gender roles can downwardly bias the evaluations of women's competence and suitability for professional and high-authority positions (Blair-Loy 2009; Correll, Benard, and Paik 2007; Cotter, Hermsen, and Vanneman 2011; Rivera 2017). These factors can also create barriers to women's employment gains in programming-intensive occupations.

Gender Differences in Wage Return Trends

The social and organizational perspectives also offer key insights into the gender-uneven trends in wage returns. First, the labor queue perspective discussed previously suggests that women may face greater barriers in advancing their careers and obtaining higher wages in programming-intensive fields because they are likely to be placed at lower ranks in the queues for higher-paying jobs. Essentialist beliefs about women's lack of intrinsic interest and abilities in tasks that are stereotyped as "male tasks," such as programming, coding, and symbolic calculations, often advantage men relative to women by encouraging gender-biased performance expectations and evaluations (Castilla 2008; Petersen and Saporta 2004; Reskin 2000; Ridgeway and Correll 2004). This in turn leads to structural disadvantages for women in promotion opportunities, likelihood of external offers, and bargaining power in salary raises in technology companies (Ahuja 2002; Tai and Sims 2005; Truman and Baroudi 1994). Women who perform programming-related tasks may still be disproportionately concentrated in relatively lower-paying positions, whereas their male counterparts on average enjoy increasingly greater wage returns associated with the rise of programming.

Second, the rise of programming in the labor market can also perpetuate preexisting, gender-biased organizational processes and thereby deepen long-standing gender inequality in the workplace. These organizational processes may take the forms of the exclusion of women from informal networks, intentional and unintentional same-sex preferences, and opportunity hoarding by already advantaged male workers (Acker 1990; Gorman 2005; Gorman and Kmec 2009; Phillips 2005; Reskin and McBrier 2000; Tomaskovic-Devey 1993). Women in professional occupations also face harsher judgment and must meet a higher standard of performance than men when it comes to the evaluation of their work (Foschi, Lai, and Sigerson 1994; Gorman 2006; Heilman 2002). These processes may be particularly salient in programming-related jobs, which arguably involve substantive training and learning and are male dominated. Over time, these organizational processes may isolate women from productivity-enhancing resources (e.g., opportunities for training and engaging with core projects) and information (e.g., opportunities for networking or job vacancies), thereby keeping them from ascending to higher-earning positions.

Third, family responsibilities, as discussed previously, can also impede women's career advancement in programming-intensive occupations. Women in the information technology (IT) industry perceive family responsibilities as hindering their advancement opportunities and being incompatible with their work schedules (Armstrong et al. 2007; Liu and Wilson 2001). Work–family constraints and conflict also lead to women's loss of productivity-enhancing work experience when they do not work as many hours as their male counterparts (Cha and Weeden 2014). To the extent that advancement in programming-intensive occupations requires substantial investment in human capital throughout the career, the growing wage returns to programming-intensive jobs may disproportionately benefit men relative to women.

Educational Differences

The labor market prospects of workers with and without a college degree may be affected differently by technological change. Workers with different education levels hold different types of jobs and face different opportunities and constraints in their careers at times of technological change. On average, programming-intensive jobs tend to require high-level skills, such as mathematical knowledge, logical and structured thinking, and knowledge of one or several programming languages, all of which are typically acquired in college or graduate school. Therefore, in times of technological upgrading, college graduates may enjoy greater opportunities for promotion and wage growth than those without a college degree. In addition, as discussed previously, the gender differentials in the educational pipeline can also contribute to the uneven impact of computer technology by gender, particularly among those with a college degree. This means that when the labor market demand for programming skills and tasks increases, the gender gap in the employment share and wages in programming-intensive occupations may grow at a particularly high rate among college graduates.

Empirical Analysis

Step 1: Constructing Occupational Measures

The need for occupation-level skill and task measures. To capture the skill and task content of the labor market, previous research has relied on individual-level measures such as whether a person uses a computer at work (Krueger 1993), aggregate-level measures such as the industry-level computer-related expenditure (Kristal and Cohen 2015), industry-level measures such as computer utilization (Autor et al. 1998) and investment in computing machinery (Autor et al. 1998; Berndt and Morrison 1995), and firm-level measures of IT innovation adoption (Bresnahan, Brynjolfsson, and Hitt 2002; King, Reichelt, and Huffman 2017). Departing from this prior literature, we propose a new type of measure taken on the level of 386 detailed occupations. We argue that there are two major advantages of occupation-level measures.

First, an important feature of recent advances in computer technology is that, rather than being contained within the IT industry, the increasing demand for programming is much more pervasive, extending to occupations outside of the IT industry. Hence, industry-based measures may miss a substantial portion of the labor market that has been affected by the rise of programming. In addition, measures constructed on the level of disaggregated occupational categories allow us to capture gender-specific changes within industries. Although individual-level skill or task measures may be an even better alternative, we know of no large-scale, repeated, individual-level measures available over the past two decades.

Second, we echo previous stratification literature in emphasizing that detailed occupations present appropriate "microclasses" (Grusky and Sørensen 1998; Weeden 2002; Weeden and Grusky 2005). Detailed occupational groups also represent divisions of labor within which relatively similar tasks are performed by practitioners (Acemoglu and Autor 2011; Deming 2017; Liu and Grusky 2013). Compared to macroclass categories or industry groups, detailed occupational categories are more powerful in explaining individual and group differences in earnings and life chances (Levanon and Grusky 2016; Liu and Grusky 2013; Mouw and Kalleberg 2010; Weeden and Grusky 2005). Although job characteristics may still vary within occupations (Tomaskovic-Devey 1993; Xie, Killewald, and Near 2016), the inclusion of a rich set of individual-level covariates that are strong predictors of labor market outcomes in our empirical analysis will further help minimize these variations. Specifically, our analysis uses the harmonized occupation schemes based on the 1990 census occupation classification scheme.[2] Data from years 2003 onward were back coded to the harmonized 1990 classification scheme using the strategy developed by the Bureau of Labor Statistics (BLS) for long-term comparisons (Meyer and Osborne 2005). For example, to back code occupations from 2003 onward to the harmonized 1990 scheme, the BLS has relied on the fact that from 2000 to 2002, many Current Population Survey (CPS) records were dual coded into the 1990 census category system as well as the 2000 census system, which provides additional information on the splitting of occupational categories in the census classification scheme over time (Meyer and Osborne 2005). Despite the various checks that were incorporated in the BLS's harmonization procedure, we note that this procedure has its own limitations, which we will come back to in the "Robustness Checks and Limitations" section.

We next discuss how we construct our occupation-level measures. Importantly, in contrast to prior work that has mainly focused on the trends in overall computer usage regardless of what tasks are being performed on a computer, we differentiate between two distinct dimensions: programming intensity and general computer usage intensity.

Task-intensity measures from the CPS-CIU data. The task-intensity measures are based on data from the Computer and Internet Use (CIU) supplement of the CPS, accessed from the Integrated Public Use Microdata Series (IPUMS) database (Flood et al. 2018). In the CIU supplement, employed persons aged 15 and older were first asked a question about whether they directly use computers at work. Those who responded "yes" were then asked if they performed one or more of these tasks: programming, spreadsheets or databases, word processing or desktop publishing, scheduling, emails, and graphics. We use the CIU data from 1997, 2001, and 2003.[3] Based on respondents' answers, we calculate the share of practitioners in an occupation who perform each of these tasks at work. This calculation is conducted for the three years separately and then pooled together. Because results are similar either way, we focus on the latter, as it provides a larger sample size for each occupation. All analyses are weighted to reflect the population distribution.

We next identify the latent task-intensity dimensions based on the six occupation-level variables calculated above using exploratory factor analysis. The factor analysis identified a total of three factors. The factor loadings after orthogonal rotation are presented in Appendix C in the online supplement. Factor 1, on which the occupational share of workers who perform programming tasks at work is loaded to a high degree (0.703) and the other variables are loaded to a substantially smaller degree, corresponds to programming intensity. Factor 2 corresponds to general computer usage intensity, as the occupational shares of workers who perform tasks related to spreadsheets or databases, word processing or desktop publishing, scheduling, and emailing are loaded to a high degree (ranged between 0.756 and 0.873) on this factor. Factor 3 corresponds to graphics intensity, as the occupational share of workers who perform tasks related to graphics has a dominating loading (0.858). Finally, we use the predicted factor scores resulting from the orthogonal rotation as our task-based measures for occupation-level programming intensity, general computer usage intensity, and graphics intensity, respectively.

*Additional data source: skill-intensity measures from O*NET data*. Next, we introduce the Occupational Information Network (O*NET) as an additional data source to construct occupation-level skill-intensity measures. In the interest of space, we leave the details of the procedure for setting up the O*NET measures in Appendix B in the online supplement. The final skill measures used in our analysis include nine specific skill factors grouped into four categories: cognitive (verbal, quantitative, and analytic), creative, technical (programming, general computer, science and engineering, and technical miscellaneous), and social (managerial and care work).

The addition of the O*NET data serves two important purposes. First, because the O*NET data also contain information on the occupational requirement for computer-related skills, they provide leverage for checking the validity of our occupation-level task-intensity measures constructed from the CPS-CIU data set above (details in the next subsection). Second, the O*NET data provide a range of additional occupation-level skill measures that are beyond computer usage per se, such as analytic, verbal, and management skills. Controlling for these additional skill measures in our models enables us to obtain the net effect of computer or programming intensity that operates independently of the other dimensions of occupational skills (Liu and Grusky 2013).

Confirming the validity of task and skill measures. Table 1 presents the 20 highest ranked nonmanual occupations in programming and general computer usage intensity by our task-based measure.[4] In terms of programming intensity, mathematicians and computer software developers rank at the top. The rest of this list contains several engineering-related occupations, jobs that involve the use of programming to help create and analyze mechanical designs, run simulations, interact with connected systems, and generate specifications for their products. In terms of general computer usage, a number of subject-specific (postsecondary) instructors rank at the top. The list also includes professional occupations whose job duties involve intensive interaction with computers but not programming, such as lawyers, accountants and auditors, and writers and authors, as well as several types of semiskilled nonmanual occupations, such as insurance adjusters, examiners, and investigators. The lists for programming and general computer usage have only two overlapping occupations: physicists and astronomers, and physics instructors (postsecondary). Overall, these descriptive results suggest that our occupation-level programming intensity and general computer usage intensity measures capture two quite distinct dimensions of occupational characteristics.

Table 1: Twenty highest ranked nonmanual occupations in programming intensity and general computer usage intensity.

Occupation	Factor Score
Panel A: Top-ranking occupations in programming intensity	
Mathematicians and mathematical scientists	6.79
Computer software developers	5.30
Computer systems analysts and computer scientists	3.86
Programmers of numerically controlled machine tools	3.81
Engineering instructors (postsecondary)	3.41
Actuaries	2.76
Statisticians	2.47
Electrical engineers	2.17
Chemistry instructors (postsecondary)	1.91
Physicists and astronomers	1.79
Aerospace engineer	1.76
Chemical engineers	1.73
Electrical and electronic (engineering) technicians	1.61
Metallurgical and materials engineers, variously phrased	1.42
Not-elsewhere-classified engineers	1.36
Sales engineers	1.21
Physics instructors (postsecondary)	1.15
Operations and systems researchers and analysts	1.13
Mechanical engineers	1.02
Mechanical engineering technicians	1.00
Panel B: Top-ranking occupations in general computer usage intensity	
History instructors (postsecondary)	3.02
Economics instructors (postsecondary)	2.51
Law instructors (postsecondary)	2.25
Lawyers	2.21
Education instructors (postsecondary)	2.19
Physicists and astronomers	2.15
Insurance underwriters	2.15
Theology instructors (postsecondary)	2.14
Sociology instructors (postsecondary)	2.14
Physical scientists, n.e.c.	2.08
Legal assistants, paralegals, legal support, etc.	2.08
Other financial specialists	1.87
Writers and authors	1.87
Physics instructors (postsecondary)	1.84
Human resources and labor relations managers	1.84
Insurance adjusters, examiners, and investigators	1.83
Accountants and auditors	1.83
Purchasing managers, agents and buyers, n.e.c.	1.79
Human resources clerks, except payroll and timekeeping	1.78
Financial managers	1.78

See Figure 1 for the data source. n.e.c., not elsewhere classified.

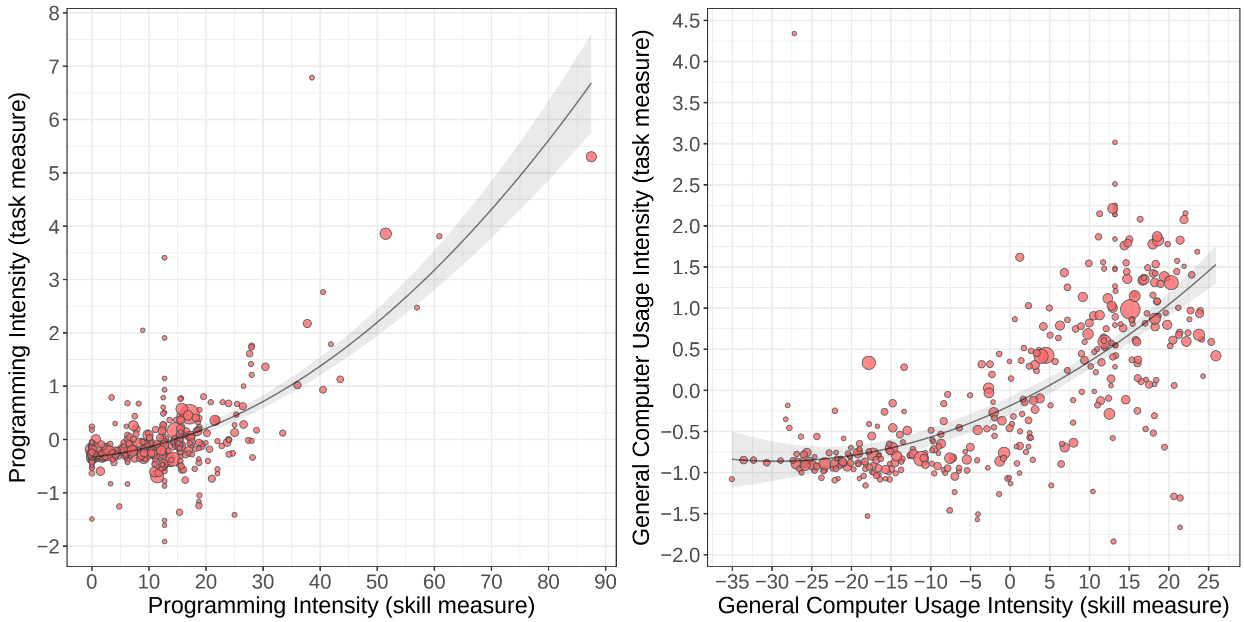

Figure 1: Occupation-level task- and skill-intensity measures of programming and general computer usage. Occupation-level task-intensity measures are based on data from the 1997, 2001, and 2003 waves of the Computer and Internet Use supplement of the CPS; occupation-level skill-intensity measures are based on data from the Occupational Information Network. Dot sizes indicate employment sizes of the corresponding occupation. The solid line and shaded area represent the quadratic fit and 95 percent confidence interval, respectively.

A second method for assessing the validity of our measures relies on the premise that, arguably, the skill and task measures pertaining to the same type of computer usage should be positively associated with each other. Hence, we next check the convergence validity of our occupation-level measures by examining their associations. Figure 1 presents scatter plots for the skill- and task-based measures of programming intensity (left panel) and those of general computer usage intensity (right panel). The fact that our task-based measures are positively associated with the corresponding skill-based measures confirms the validity of our measures.

In summary, our measurement strategies consistently identify programming and general computer usage as two distinct dimensions of computer usage at work. In all analyses below, we conducted parallel analyses using both skill- and task-based measures, which have yielded very similar findings. We thus present here findings based on the task-intensity measures only.

Step 2: Estimating Gender-Specific Trends in Occupational Composition

This step examines the trends in the programming intensity of men's and women's jobs over time. To do so, we link the measures constructed in step 1 to data from the CPS, a monthly, nationally representative survey of employment and labor markets.

We use the 1994 to 2015 waves of the merged outgoing rotation groups files. We focus on this period for two reasons. First, unlike prior computerization literature that has focused primarily on the rise of general computer usage in years prior to the mid-1990s, our study focuses on the rise of programming, which took off around the mid-1990s (Acemoglu and Restrepo 2017; Autor 2015). Second, because our occupational task intensity measures come from the CPS-CIU supplement in 1997 or later, extending the time window further back will likely require strong assumptions about the stability of task intensity over time.

Has the programming intensity and general computer usage intensity of men's and women's jobs changed over time? Our results suggest that these trends differ dramatically by gender and education. Figure 2 presents the trends in average programming intensity (upper panel) and general computer intensity (lower panel) by gender and by whether one has a college degree.[5] College-educated men hold jobs with the highest level of programming intensity, and the average factor score increases steadily over the two decades. By contrast, men without a college degree and women with a college degree have jobs with very similar levels of programming intensity on average, and both trends have remained rather stable over time. Although women with a college degree did exhibit a moderate increase from the mid-1990s to 2000 and after 2010, this increase is much smaller than that for college-educated men.

The lower panel of Figure 2, showing the trend in general computer intensity, provides a placebo benchmark for the observed trends in programming intensity above. That is, if the rise of programming intensity is simply a natural result of the continued proliferation of computer technology in the labor market in general rather than anything related to programming per se, then we would expect to see a similarly strong increase in general computer intensity among college-educated men. However, our results indicate exactly the opposite. Among both college graduates and those without a college degree, women's occupations tend to have higher general computer intensity scores, although the gender gap is larger among those without a college degree. There is no growth in general computer intensity across the labor market, and in fact, the average factor scores have declined gradually among college-educated men and women without a college degree. These findings suggest that the rise in programming reflects a unique dimension of technological changes that took place in the post-1990s period.

Step 3: Estimating Gender-Specific Trends in Wage Returns to Programming Intensity

We next estimate the wage returns to occupational programming intensity. The outcome variable in this set of analyses is hourly wage. For workers paid by the hour, this variable is based on their reported hourly wage; for workers who are not paid by the hour, hourly wage is calculated as the amount usually earned per week at the respondent's current job divided by the usual number of hours worked per week. CPS includes several different measures for weekly work hours, and we use the usual weekly hours corresponding to the respondent's current job because

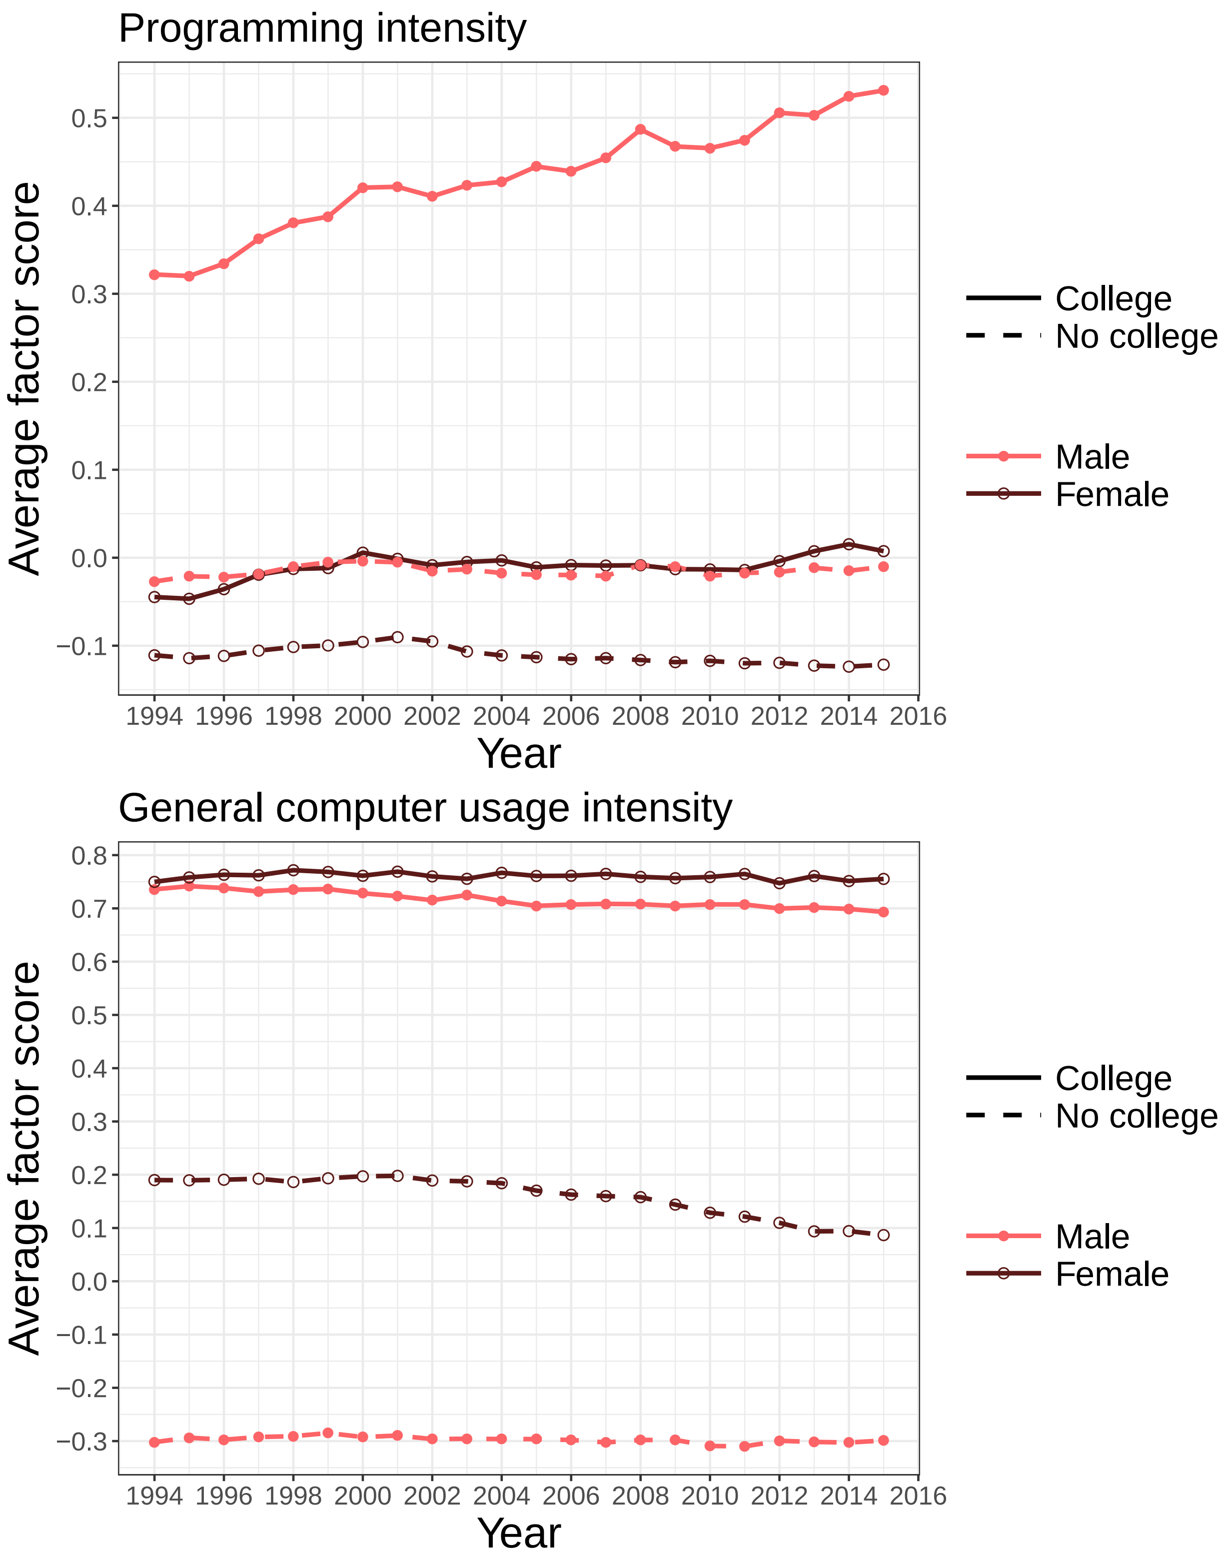

Figure 2: Gender-specific trends in average factor scores of programming intensity and general computer usage intensity. See Figure 1 for the data source.

this specification is most consistent with the weekly earnings variable above. Self-employed individuals are excluded from the analysis. We adjust for inflation by converting wages to 1999 dollars. As per standard practice (e.g., Card and DiNardo 2002; Mouw and Kalleberg 2010), top-coded wages were assigned 1.5 times the top-coded value. Hourly wage captures the economic returns to an hour of labor and therefore avoids conflating the gender difference in hourly wage and work hours.[6] Our sample includes workers with both full- and part-time employment to avoid substantial selection bias into being a full-time worker (Card and DiNardo 2002), but we include a full-time/part-time dummy in the wage regression models to account for their wage differentials.

The baseline model (M1) predicts log hourly wage (Log(W)) as:

$$\begin{aligned} Log(W) = \beta_0 &+ \beta_{prog} \cdot \delta^{prog} \\ &+ \text{ other computer usage controls} \\ &+ \text{ demographic controls,} \end{aligned} \tag{1}$$

where δ^{prog} is the standardized factor score for programming intensity. That is, one unit change in this standard programming intensity variable represents one standard deviation change of the raw factor score in the population. Other computer usage controls include standardized factor scores for general computer intensity and graphic intensity. These controls help ensure that the estimated wage return to programming intensity is not an artifact of the possibility that programming-intensive occupations also tend to rely on other computer-related tasks. Demographic controls include age, age squared, and dummies for race (white, black, and Asian).

After the baseline model, we estimate three additional models: M2 adds to M1 controls for educational attainment (less than high school, high school graduate, some college, bachelor's degree, master's degree, and professional or doctoral degree). In our additional analysis, we interacted education with programming-intensity measures, and the results indicate that the wage effect of programming-intensity does not vary significantly by education levels, and we therefore include these educational attainment dummies only additively in our models. M3 further adds controls for employment characteristics (union status, full-time/part-time status, and industrial categories) to M2. Finally, M4 adds controls for occupation-level skill importance scores (except for the scores for general computer knowledge and programming skills because of their high degree of overlap with our key task-intensity measures) as described in step 1. Descriptive statistics for these control variables are presented in Appendix D in the online supplement.

Table 2 presents the estimated coefficients from the models estimated for men and women separately. We pool the years into two periods (1994 to 2004 and 2005 to 2015) for the sake of presentation. As the dependent variable is log transformed, the effects can be interpreted in percentage terms. In the baseline model (M1), among men, one standard deviation increase in programming intensity is associated with 5.85-percent and 6.52-percent increase in hourly wage during the 1994 to 2004 and 2005 to 2015 periods, respectively. Adding educational controls (M2) and employment characteristics (M3) reduces the wage returns to programming intensity

Table 2: Selected coefficients in regression models predicting log hourly wage by gender and period.

	1994–2004				2005–2015			
	(1) M1	(2) M2	(3) M3	(4) M4	(5) M1	(6) M2	(7) M3	(8) M4
Panel A: Men								
Programming	0.0585*	0.0544*	0.0510*	0.0253*	0.0652*	0.0565*	0.0553*	0.0269*
	(92.59)	(87.21)	(79.52)	(30.62)	(102.48)	(90.41)	(85.72)	(29.20)
BIC	976,878.1	936,334.4	898,216.4	883,109.4	986,381.3	945,330.6	916,472.7	900,494.6
R^2	0.231	0.279	0.321	0.334	0.236	0.286	0.320	0.331
Panel B: Women								
Programming	0.0735*	0.0743*	0.0760*	0.0542*	0.0763*	0.0735*	0.0769*	0.0527*
	(87.47)	(89.74)	(87.72)	(50.67)	(68.33)	(67.65)	(68.66)	(38.35)
BIC	943,564.6	865,650.8	845,248.7	828,717.7	976,329.6	906,750.2	894,457.3	878,504.7
R^2	0.187	0.284	0.307	0.321	0.186	0.277	0.292	0.305
Other computer usage controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education controls	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Employment/industry controls	No	No	Yes	Yes	No	No	Yes	Yes
Skill requirement controls	No	No	No	Yes	No	No	No	Yes

t statistics in parentheses. Data come from the 1994 to 2015 waves of the merged outgoing rotation group files of the Current Population Survey. Occupation-level task-intensity measures are based on data from the Computer and Internet Use supplement of the Current Population Survey. Hourly wages are adjusted to 1999 dollars. All statistics are weighted to represent the population at each year.
$^*p < 0.01$

marginally. Finally, in M4, the wage returns reduce to 2.53 percent in 1994 to 2004 and 2.69 percent in 2005 to 2015, suggesting that about half of the total wage returns to programming-intensity among men are explained by the fact that programming-intensive occupations also tend to be associated with high-paying skills. The wage payoff to programming intensity is statistically significant at the 0.001 level for both genders. Among women, in the model with the full set of controls (M4), one standard deviation increase in the individual's occupation's programming intensity is associated with 5.42-percent and 5.27-percent wage returns in the two periods, respectively. This suggests that although women experience a wage disadvantage relative to men on average, compared to their same-gender counterparts, working in an occupation with higher programming intensity is associated with a greater wage payoff for women than for men.

Note also that M4 fares better than the other models in terms of goodness of fit, which is assessed here using the R^2 and BIC statistics shown in Table 2. A higher R^2 or lower Bayesian information criterion (BIC) score indicates a better model fit. The two statistics consistently favor M4 for both genders in our two periods.[7] Therefore, all our following analyses will be based on M4.

One way to further inspect the regression results is to compare the relative sizes of the wage premium on occupation-level task intensity and the educational premium. Table 3 presents these results. For simplicity, we refer to the wage returns to one standard deviation of programming or general computer intensity as the technological premium, the wage returns to higher educational attainment relative to a high school diploma as the educational premium, and the ratio between the two premiums as the relative wage premium. For example, in Table 3, $\beta^{programming}/\beta^{BA}$

Table 3: Comparison of technological and educational wage premiums.

	1994–2004	2005–2015
Panel A: Men		
Technological and educational wage premiums		
$\beta^{programming}$	0.025	0.027
$\beta^{general}$	0.070	0.055
β^{BA}	0.254	0.290
β^{MA}	0.347	0.384
Relative wage premium (technological premium versus educational premium)		
$\beta^{programming} / \beta^{BA}$	9.96%	9.28%
$\beta^{programming} / \beta^{MA}$	7.29%	7.01%
$\beta^{general} / \beta^{BA}$	27.56%	19.00%
$\beta^{general} / \beta^{MA}$	20.17%	14.35%
N	632,187	600,214
Panel B: Women		
Technological and educational wage premiums		
$\beta^{programming}$	0.054	0.053
$\beta^{general}$	0.165	0.174
β^{BA}	0.340	0.347
β^{MA}	0.459	0.454
Relative wage premium (technological premium versus educational premium)		
$\beta^{programming} / \beta^{BA}$	15.94%	15.19%
$\beta^{programming} / \beta^{MA}$	11.81%	11.61%
$\beta^{general} / \beta^{BA}$	48.53%	50.14%
$\beta^{general} / \beta^{MA}$	35.95%	38.33%
N	613,453	584,841

See Table 2 for the data source.

represents the relative wage premium between programming intensity and a bachelor's degree. Note that these wage premiums are obtained from the model (M4) that controls for a rich set of demographic, employment, and occupational characteristics. Among men, the wage premium associated with programming intensity is as large as about 10 percent of the college wage premium and about 7 percent of a master degree's wage premium. Although the educational wage premium has increased across the two periods, the relative wage premium of programming intensity remained largely unchanged across the two periods, suggesting that the wage returns to programming intensity have actually kept pace with the wage returns to a bachelor's or master's degree. The wage premium of general computer intensity relative to the educational wage premium, however, decreased by about one-third from 1994 through 2004 to 2005 through 2015, suggesting that among men, the returns to general computer intensity have not kept pace with the educational wage premium.

To examine the trends over time, we estimate the full model, M4, for every year from 1994 to 2015. We report the estimated trends in the wage returns to education, skill measures, and race in Appendix H in the online supplement, and focus our discussion here on the wage returns to programming and general computer usage presented in Figure 3. The lines and shaded areas represent the Loess-smoothed curves of the trends and their 95% confidence intervals.[8] For both genders, the wage returns to general computer usage first declined moderately and then increased again, with changes being larger among men. Although direct examination of this trend is beyond the scope of this study, we note that the trends in general computer usage intensity may reflect the business cycle and that the rise after 2010 reflects some recovery. The temporal variations in the wage returns to general computer usage are quite mild among women, which may be due to the fact that men's labor market outcomes are generally more responsive to business cycles. But overall, the magnitude of these temporal changes in the wage returns to general computer usage is much smaller compared to the gender-specific trends in the wage returns to programming intensity. As the upper panel suggests, wage returns to programming intensity have remained almost unchanged among female employees over this period. By contrast, the most notable change in this figure is the increase in the wage returns to programming intensity among men starting around 2006. In 2015, the wage premium associated with programming intensity among men increased by about 60 percent compared to its level in 1994. In sum, this part of our analysis reveals a gender difference in the "price" trends: Men, but not women, have enjoyed the growing economic returns associated with programming intensity, particularly in the most recent decade.

Step 4: Composition and Price Effects on the Stalled Convergence of the Gender Wage Gap

Finally, to what extent does the rise of programming-intensive occupations affect the trend in the gender wage gap? To evaluate the relative contributions of the composition and price effects, we predict three counterfactual trajectories of the gender wage gap from 1994 to 2015 and compare them to the observed trend in the gender wage gap. The first counterfactual "turns off" the price effect by fixing the wage returns to programming intensity at the 1994 level while allowing the composition to vary. The second counterfactual turns off the composition effect by fixing the labor market composition in terms of occupational programming intensity while allowing the wage returns to vary. The third counterfactual turns off both composition and price effects by fixing the labor market composition and price at the 1994 level. We do this exercise separately for those with a college degree and those without. Appendix J in the online supplement provides the technical details for calculating these counterfactuals. Here, we use these counterfactual predictions not to draw a causal claim but for the purpose of discerning the relative contributions of composition and price effects to the overall trends in the gender wage gap. These counterfactual trajectories are presented in panel A of Figure 4.

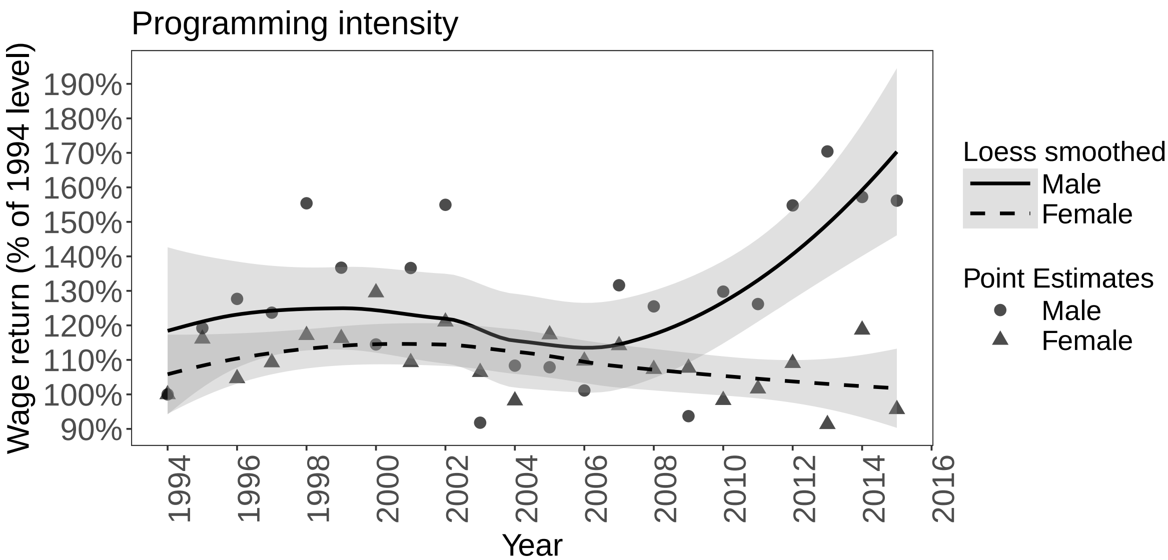

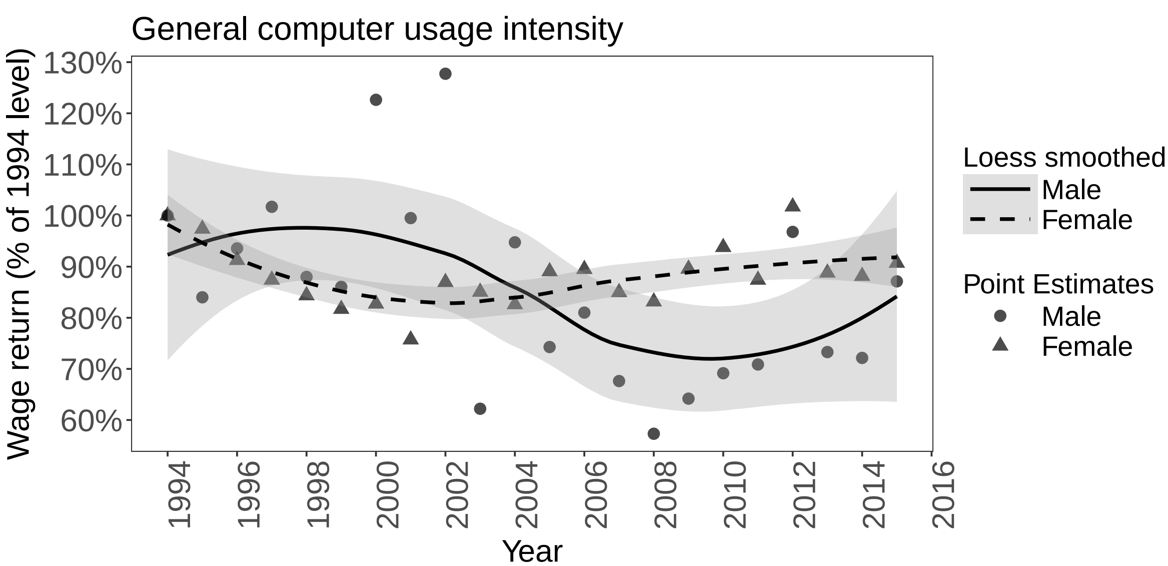

Figure 3: Gender-specific trends in wage returns (as percentage of 1994 level) to programming intensity and general computer usage intensity. See Table 2 for the data source. Solid and dashed lines represent the Loess-smoothed curve with a span of 0.5, and shaded areas represent 95 percent confidence intervals. Alternative span widths yield similar findings.

Then, in panel B of Figure 4, we present the trajectories of the composition and price effects by taking the differences between the observed gender wage gap and the predicted counterfactuals. The slope of these trajectories indicates whether these effects facilitated or impeded the closing of the gender wage gap. An upwardly sloping trajectory indicates that the effect has impeded the narrowing of the gender

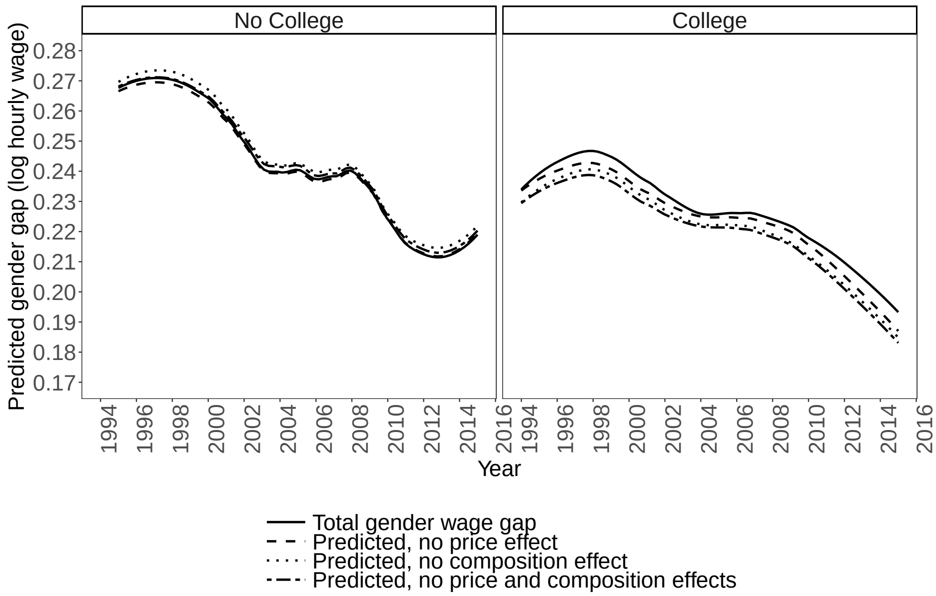

(a) Predicted gender wage gap.

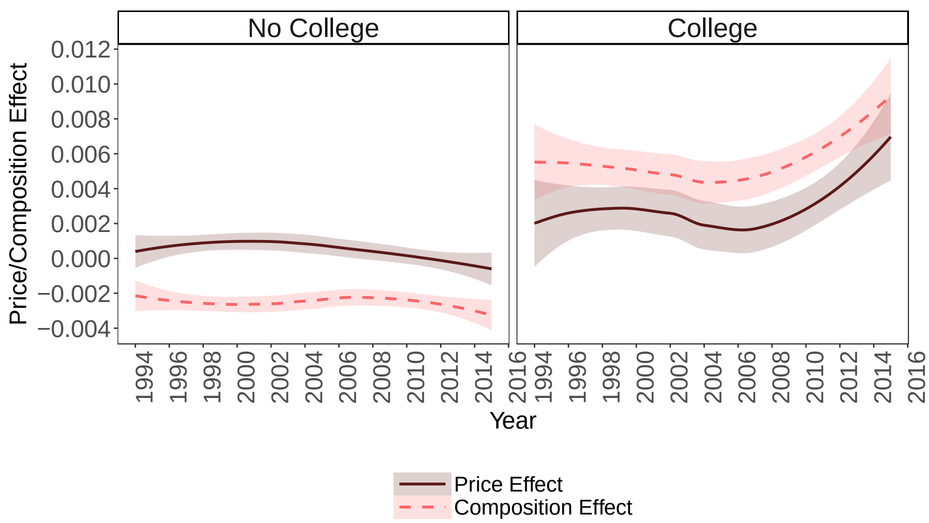

(b) Composition and price effects.

Figure 4: Counterfactual simulation results for composition and price effects of programming intensity on the gender wage gap. See Table 2 for the data source. All curves are Loess smoothed with a span 0.5. Alternative span widths yield similar findings.

gap and vice versa.[9] The trajectories are flat among those without a college degree, suggesting that neither the composition nor the price effect of programming intensity has affected changes in the gender wage gap. By contrast, among those with a college degree, both composition and price effects have increased over time, which suggest that they both impeded the convergence of the gender wage gap over time. These effects are particularly large in the most recent decade.

Finally, to quantify the contributions of the composition and price effects to the stalled convergence of the gender wage gap over this period, we calculate the changes in these effects from 1994 to 2015 and present the results in Table 4. For both education groups, the gender gap in log hourly wage has declined over time, and the decline is greater among those without a college degree (–0.0602 versus –0.0380 on log wage scale). The last column of the table expresses the change in the gender wage gap explained by the price and composition effects as percentages of the total change in the gender wage gap. A positive percentage means that the price or composition effect has facilitated the movement towards gender wage equality, and a negative percentage means the effect has impeded the movement towards gender wage equality. Note that the total effect does not necessarily equal the sum of composition and price effects because the two effects are multiplicative rather than additive and therefore have an overlapping portion (see Appendix J in the online supplement for a technical demonstration). Among those without a college degree, the composition, price, and total effects are positive but very small (1.51 percent, 1.72 percent, and –0.34 percent, respectively), indicating that the rise of programming-intensive occupations has had little impact on the gender wage gap. By contrast, among college graduates, both effects have substantially hindered the convergence in the gender wage gap. Without the price and composition effects, the gender wage gap would have converged by an additional 11.76 percent and 14.18 percent, respectively. The total contribution of the two mechanisms comes to 14.70 percent of the observed convergence in the gender wage gap. That is, the absence of these two effects would have reduced the gender wage gap over the past two decades by an additional 14.70 percent.

Robustness Checks and Limitations

Our findings above are robust to alternative model specifications in which we (1) use skill-intensity instead of task-intensity measures for programming and general computer usage, (2) restrict the sample to the nonmanual sector, and (3) switch to American Community Survey data (but with a shorter time period that was available, see Appendix K in the online supplement). To ensure that our results are not driven by a small number of the most programming-intensive occupations, we also reestimated all the models with the top 10 most programming-intensive occupations removed from the sample, which generated consistent findings (see Appendix L in the online supplement). Finally, we explored the variations in the gender-specific trends within seven major industrial categories: manufacturing, wholesale/retail, personal service, business service, finance, professional, and public administration. The results indicate that the divergence in programming intensity between jobs held by college-educated men and women is evident across various industries, with the most dramatic divergence occurring in the finance industry (see Appendix M in the online supplement). Overall, this suggests that the

Table 4: Contributions of the composition and price effect of programming intensity to the convergence of the gender wage gap from 1994 to 2015.

	1994	2015	Change	% of Change in Gender Gap
No college				
Gender gap in log wage	0.2818	0.2216	−0.0602	
Composition effect	−0.0021	−0.0030	−0.0009	1.51%
Price effect	0.0000	−0.0010	−0.0010	1.72%
Total effect	−0.0021	−0.0019	0.0002	−0.34%
College				
Gender gap in log wage	0.2373	0.1993	−0.0380	
Composition effect	0.0042	0.0086	0.0045	−11.76%
Price effect	0.0000	0.0054	0.0054	−14.18%
Total effect	0.0042	0.0098	0.0056	−14.70%

labor-market–wide trend in gender differences shown in Figure 2 also holds within industries.

Our findings should be interpreted in the context of several limitations. First, our analysis employs time-invariant occupation-level measures, a decision driven by the availability of computer task measures and the need to ensure sufficient data points within detailed occupations. In doing so, we might have ignored some changes in the skill or task content of occupations. For instance, studies show that software developers, computer support specialists, computer and information systems managers, and computer systems analysts have grown into the largest detailed occupations among IT workers in 2014, and these IT jobs have also increased in terms of the specificity of their tasks (Beckhusen 2016). For example, the category of "network systems and data communications analysts" in the 2000 Standard Occupational Classification split into four subcategories in its 2010 version: information security analysts, web developers, network and computer systems administrators, computer network architects, and computer network support specialists (BLS 2010). We note that it is likely that some of these changes impacted men and women differently. For example, within the same detailed occupation that involves various types of computer-related tasks, men may still be more likely to specialize in more programming-related tasks, whereas women are more likely to specialize in general computer usage. In this case, we might have underestimated the true effect of programming on the gender wage gap by failing to account for within-occupation gendered division of labor. Examining trends in gender inequality and segregation within the level of detailed occupations is beyond the scope of the current study and should be further examined by future work.

Second, our analysis focuses on occupation and wage outcomes among those who currently have a job. But the rise of programming may have also affected the likelihood of finding, changing, or losing employment. This trend may have also affected the long-term career trajectories through which individuals acquire skills, accumulate work experience, and achieve wage growth. To better understand the impact of the rise of programming on the full range of employment outcomes, future research is needed to examine how programming affects the dynamic transitions in the labor market.

Third, our analysis uses a harmonized occupational coding scheme, which assigns universal conversion factors for both genders as opposed to gender-specific weights, but previous work has pointed out that the back coding may mask changes in the gender ratios of occupations from 2003 onward, when the back coding was applied (Weeden 1998). Although it is beyond the scope of this current article to examine these gender-specific patterns in occupational coding schemes, as a robustness check, we examine visually whether the trends in composition and price changes remain robust within periods when the same census occupational classification schemes were used by the CPS. The results, shown in Figures N1 and N2 in Appendix N in the online supplement, suggest that the findings are robust within periods in which the same occupational classification scheme is used.

Lastly, as discussed earlier, gender differentials in the labor market may well have occurred on the supply side prior to entry into the labor market, especially in the educational system. Gender differences in college majors as well as the school-to-work pipeline may have a strong role in explaining differences in employment and wage outcomes between men and women. Hence, our findings of a widening gender gap among college graduates open avenues for future research as to whether these trends are attributable to the gender differences in college majors.

Conclusion

This study started with the proposition that whereas the pre–mid-1990s waves of computerization have favored the productivity of office jobs, which women tend to be employed in, the changes in computer technology after the mid-1990s, which have led to an increasing reliance on programming in the labor market, may have favored men over women and thus have impeded the movement toward gender equality. This proposition is supported by our empirical investigation, in which we documented a strong relationship between the rise of programming-intensive occupations and the stalled convergence in the gender wage gap among those with a college degree in the recent two decades. This relationship stems from two effects: a composition effect in which college-educated men, but not their female counterparts, are increasingly employed in occupations with high programming intensity; and a price effect in which college-educated men, but not their female counterparts, have enjoyed a growing wage payoff to working in programming-intensive occupations. If there had not been any change in the gender-specific compositions and wage returns to programming intensity, the gender wage gap would have converged by an additional 14.70%.

To put it simply, as neutral as it may appear to be, the impact of technological change turns out to be highly gendered. We showed that the recent rise of programming intensity has negatively affected gender equality, and this impact has been pervasive, spreading across a wide range of industries. Although our empirical analyses clearly demonstrate these uneven trends, the mechanisms behind this relationship may be deep rooted, multifaceted, and far from adequately understood. These complex mechanisms may include gender-biased status beliefs in the workplace, male-dominated occupational culture, organizational processes that exclude women from resources and opportunities, gender-biased performance evaluations, and persistent gendered division of labor in the family. Furthermore, such gender imbalance may also influence other domains of personal and social life. For example, the labor market context may affect men's and women's gender ideology at work and in the family (Cha and Thébaud 2009; Gerson 2009; Mason and Lu 1988). Hence, the social and cultural mechanisms that have contributed to the gendered impact of new technology may themselves be aggravated by this technological change, thus perpetuating the vicious circle. The specific mechanisms underlying these empirical trends ought to be explored in future research.

On the broader theoretical level, our findings also indicate the need to go beyond a market-based perspective for understanding the impact of technological change. Unlike the bulk of SBTC literature, which primarily focuses on the supply and demand of skills and tasks regardless of individuals' gender, we show that, conditional on a host of demographic, education, employment, and skill characteristics, the rise of programming turns out to have affected men and women differently. This finding calls attention to the social and institutional perspective on labor market inequality, which suggests that technological change interacts with beliefs and practices in existing social institutions and deepens the gendered division of labor in the workplace. To further substantiate this point, we await future research to construct data and measures on skills used and tasks performed within occupations, such as those at the job or organizational level.

Understanding the problem is certainly not sufficient for solving it, but our findings lead us to propose some potentially promising directions. On the supply side, policies and programs that facilitate women's participation in STEM fields of study, particularly fields related to computer science and data analysis, as well as their transition into programming-related occupations upon completing their education can help reduce some of the barriers that college-educated women face in entering and staying in programming-intensive occupations. On the demand side, strategies for counteracting the existing gender institutions may lie in the microlevel processes in organizations and firms (Acker 1990; Reskin 2000). For instance, a recent study suggests that working with organizational actors to develop interventions that produce small, measurable wins may be a promising way to reduce gender bias in technology companies (Correll 2017). There is much scope for future research to seriously engage with these complexities of gender inequality and continue exploring potential interventions for reducing gender inequality in the workplace.

Notes

1 A descriptive overview of the trend in computerization and gender inequality is presented in Appendix A in the online supplement.

2 We use the OCC1990 variable from the IPUMS database. This variable is based on the 1990 census occupation classification scheme from 1994 to 2002, on the 2000 scheme from 2003 to 2010, and on the 2010 scheme from 2011 to 2015.

3 CIU included questions on programming in 1989 and 1993. However, because of a coding error that was documented on the IPUMS website (see the CIWPROG variable description), these years cannot be used.

4 See Appendix C in the online supplement for additional details on these measures.

5 The same statistics but expressed as relative to the 1994 level are presented in Figure E1 in the online supplement.

6 Rodgers, Brown, and Duncan (1993) noted that self-reported weekly work hours may vary from time to time and may not always conform to actual work hours. We acknowledge that the hourly wage variable may be noisy because of both the temporal instability and the measurement errors in measured weekly hours. Another strategy adopted by previous work is to focus on weekly earnings instead of hourly wages as the outcome variable (e.g., Weeden 2002). Our current study, however, chooses to focus on hourly wage as the key outcome variable to avoid the systematic differences in weekly work hours by gender and occupational attributes.

7 Our models rely on a set of occupation-level task- and skill-intensity measures to capture the occupational differences in wages. An alternative strategy is to include occupation-specific fixed effects to absorb the maximum amount of between-occupation variations. Appendix F in the online supplement addresses the extent to which our skill/task model explains the total between-occupation variations as measured in the fixed-effect model.

8 Alternative time specifications, such as the linear and quadratic forms, generate consistent results. The unsmoothed trajectories, which suggest a consistent pattern, are presented in Appendix G in the online supplement.

9 The unsmoothed trajectories, which also suggest a consistent pattern, are shown in Appendix I in the online supplement.

References

Abbott, Andrew. 1988. *The System of Professions: An Essay on the Division of Expert Labor*. Chicago, IL: University of Chicago Press. https://doi.org/10.7208/chicago/9780226189666.001.0001.

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." *Handbook of Labor Economics* 4:1043–171. https://doi.org/10.1016/S0169-7218(11)02410-5.

Acemoglu, Daron, and Pascual Restrepo. 2017. "Robots and Jobs: Evidence from US Labor Markets." Working Paper No. 23285, National Bureau of Economic Research, Cambridge, MA.

Acker, Joan. 1990. "Hierarchies, Jobs, Bodies: A Theory of Gendered Organizations." *Gender & Society* 4:139–58. https://doi.org/10.1177/089124390004002002.

Ahuja, Manju K. 2002. "Women in the Information Technology Profession: A Literature Review, Synthesis and Research Agenda." *European Journal of Information Systems* 11:20–34. https://doi.org/10.1057/palgrave.ejis.3000417.

Alfrey, Lauren, and France Winddance Twine. 2017. "Gender-fluid Geek Girls: Negotiating Inequality Regimes in the Tech Industry." *Gender & Society* 31:28–50. https://doi.org/10.1177/0891243216680590.

Armstrong, Deborah J., Cynthia K. Riemenschneider, Myria W. Allen, and Margaret F. Reid. 2007. "Advancement, Voluntary Turnover and Women in IT: A Cognitive Study of Work–Family Conflict." *Information & Management* 44:142–53. https://doi.org/10.1016/j.im.2006.11.005.

Autor, David H. 2013. "The 'Task Approach' to Labor Markets: An Overview." *Journal for Labour Market Research* 46:185–99. https://doi.org/10.1007/s12651-013-0128-z. http://link.springer.com/10.1007/s12651-013-0128-z.

Autor, David H. 2015. "Why Are There Still so Many Jobs? The History and Future of Workplace Automation." *The Journal of Economic Perspectives* 29:3–30. https://doi.org/10.1257/jep.29.3.3.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *The Quarterly Journal of Economics* 113:1169–213. https://doi.org/10.1162/003355398555874.

Becker, Gary S. 1957. *The Economics of Discrimination*. Chicago, IL: University of Chicago Press.

Beckhusen, Julia. 2016. *Occupations in Information Technology*. Washington, DC: US Department of Commerce, Economics and Statistics Administration, US Census Bureau.

Berndt, Ernst R., and Catherine J. Morrison. 1995. "High-Tech Capital Formation and Economic Performance in US Manufacturing Industries: An Exploratory Analysis." *Journal of Econometrics* 65:9–43. https://doi.org/10.1016/0304-4076(94)01596-R.

Bianchi, Suzanne M., Liana C. Sayer, Melissa A. Milkie, and John P. Robinson. 2012. "Housework: Who Did, Does or Will Do it, and How Much Does It Matter?" *Social Forces* 91:55–63. https://doi.org/10.1093/sf/sos120.

Black, Sandra E., and Alexandra Spitz-Oener. 2010. "Explaining Women's Success: Technological Change and the Skill Content of Women's Work." *The Review of Economics and Statistics* 92:187–94. https://doi.org/10.1162/rest.2009.11761.

Blair-Loy, Mary. 2009. *Competing Devotions: Career and Family among Women Executives*. Cambridge, MA: Harvard University Press.

Blau, Francine D., and Lawrence M. Kahn. 2000. "Gender Differences in Pay." Working Paper No. 7732, National Bureau of Economic Research, Cambridge, MA.

Blau, Francine D., and Lawrence M. Kahn. 2006. "The US Gender Pay Gap in the 1990s: Slowing Convergence." *Industrial and Labor Relations Review* 60:45–66. https://doi.org/10.1177/001979390606000103.

BLS. 2010. "2010 SOC User Guide." U.S. Bureau of Labor Statistics. https://www.bls.gov/soc/soc_2010_user_guide.pdf.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *The Quarterly Journal of Economics* 117:339–76. https://doi.org/10.1162/003355302753399526.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York, NY: W. W. Norton & Company.

Burris, Beverly H. 1998. "Computerization of the Workplace." *Annual Review of Sociology* 24:141–57. https://doi.org/10.1146/annurev.soc.24.1.141.

Card, David, and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics* 20:733–83. https://doi.org/10.1086/342055.

Castilla, Emilio J. 2008. "Gender, Race, and Meritocracy in Organizational Careers." *American Journal of Sociology* 113:1479–526. https://doi.org/10.1086/588738.

Cech, Erin, Brian Rubineau, Susan Silbey, and Caroll Seron. 2011. "Professional Role Confidence and Gendered Persistence in Engineering." *American Sociological Review* 76:641–66. https://doi.org/10.1177/0003122411420815.

Cha, Youngjoo. 2010. "Reinforcing Separate Spheres: The Effect of Spousal Overwork on Men's and Women's Employment in Dual-Earner Households." *American Sociological Review* 75:303–29. https://doi.org/10.1177/0003122410365307.

Cha, Youngjoo, and Sarah Thébaud. 2009. "Labor Markets, Breadwinning, and Beliefs: How Economic Context Shapes Men's Gender Ideology." *Gender & Society* 23:215–43. https://doi.org/10.1177/0891243208330448.

Cha, Youngjoo, and Kim A. Weeden. 2014. "Overwork and the Slow Convergence in the Gender Gap in Wages." *American Sociological Review* 79:457–84. https://doi.org/10.1177/0003122414528936.

Charles, Maria, and Karen Bradley. 2009. "Indulging Our Gendered Selves? Sex Segregation by Field of Study in 44 Countries." *American Journal of Sociology* 114:924–76. https://doi.org/10.1086/595942.

Charles, Maria, and David B. Grusky. 2005. *Occupational Ghettos: The Worldwide Segregation of Women and Men*. Vol. 71. Münster, Germany: LIT Verlag.

Cheryan, Sapna, Victoria C. Plaut, Paul G. Davies, and Claude M. Steele. 2009. "Ambient Belonging: How Stereotypical Cues Impact Gender Participation in Computer Science." *Journal of Personality and Social Psychology* 97:1045–60. https://doi.org/10.1037/a0016239.

Cohoon, Joanne McGrath. 2001. "Toward Improving Female Retention in the Computer Science Major." *Communications of the ACM* 44:108–14. https://doi.org/10.1145/374308.374367.

Cooper, Joel, and Kimberlee D. Weaver. 2003. *Gender and Computers: Understanding the Digital Divide*. Abingdon, United Kingdom: Taylor & Francis.

Correll, Shelley J. 2001. "Gender and the Career Choice Process: The Role of Biased Self-assessments." *American Journal of Sociology* 106:1691–730. https://doi.org/10.1086/321299.

Correll, Shelley J. 2004. "Constraints into Preferences: Gender, Status, and Emerging Career Aspirations." *American Sociological Review* 69:93–113. https://doi.org/10.1177/000312240406900106.

Correll, Shelley J. 2017. "SWS 2016 Feminist Lecture: Reducing Gender Biases In Modern Workplaces: A Small Wins Approach to Organizational Change." *Gender & Society* 31:725–50. https://doi.org/10.1177/0891243217738518.

Correll, Shelley J., Stephen Benard, and In Paik. 2007. "Getting a Job: Is There a Motherhood Penalty?" *American Journal of Sociology* 112:1297–338. https://doi.org/10.1086/511799.

Cotter, David, Joan M. Hermsen, and Reeve Vanneman. 2011. "The End of the Gender Revolution? Gender Role Attitudes from 1977 to 2008." *American Journal of Sociology* 117:259–89. https://doi.org/10.1086/658853.

Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *The Quarterly Journal of Economics* 132:1593–640. https://doi.org/10.1093/qje/qjx022.

England, Paula. 2010. "The Gender Revolution: Uneven and Stalled." *Gender & Society* 24:149–66. https://doi.org/10.1177/0891243210361475.

England, Paula, and Su Li. 2006. "Desegregation Stalled: The Changing Gender Composition of College Majors, 1971–2002." *Gender & Society* 20:657–77. https://doi.org/10.1177/0891243206290753.

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, and J. Robert Warren. 2018. "Integrated Public Use Microdata Series, Current Population Survey: Version 6.0 [dataset]." University of Minnesota. https://doi.org/10.18128/D030.V6.0.

Foschi, Martha, Larissa Lai, and Kirsten Sigerson. 1994. "Gender and Double Standards in the Assessment of Job Applicants." *Social Psychology Quarterly* 57:326–39. https://doi.org/10.2307/2787159.

Freeman, Richard B. 2002. "The Labour Market in the New Information Economy." *Oxford Review of Economic Policy* 18:288–305. https://doi.org/10.1093/oxrep/18.3.288.

Frenkel, Karen A. 1990. "Women and Computing." *Communications of the ACM* 33:34–46. https://doi.org/10.1145/92755.92756.

Friedberg, Leora. 2003. "The Impact of Technological Change on Older Workers: Evidence from Data on Computer Use." *Industrial and Labor Relations Review* 56:511–29. https://doi.org/10.1177/001979390305600309.

Gerson, Kathleen. 2009. *The Unfinished Revolution: Coming of Age in a New Era of Gender, Work, and Family*. New York, NY: Oxford University Press.

Goldin, Claudia. 2002. "The Rising (and Then Declining) Significance of Gender." Working Paper No. 8915, National Bureau of Economic Research, Cambridge, MA.

Goldin, Claudia, and Lawrence F. Katz. 1998. "The Origins of Technology-Skill Complementarity." *The Quarterly Journal of Economics* 113:693–732. https://doi.org/10.1162/003355398555720.

Gorman, Elizabeth H. 2005. "Gender Stereotypes, Same-Gender Preferences, and Organizational Variation in the Hiring of Women: Evidence from Law Firms." *American Sociological Review* 70:702–28. https://doi.org/10.1177/000312240507000408.

Gorman, Elizabeth H. 2006. "Work Uncertainty and the Promotion of Professional Women: The Case of Law Firm Partnership." *Social Forces* 85:865–90. https://doi.org/10.1353/sof.2007.0004.

Gorman, Elizabeth H., and Julie A. Kmec. 2009. "Hierarchical Rank and Women's Organizational Mobility: Glass Ceilings in Corporate Law Firms." *American Journal of Sociology* 114:1428–74. https://doi.org/10.1086/595950.

Grusky, David B., and Jesper B. Sørensen. 1998. "Can Class Analysis Be Salvaged?" *American Journal of Sociology* 103:1187–234. https://doi.org/10.1086/231351.

Heilman, Madeline E. 2002. "Description and Prescription: How Gender Stereotypes Prevent Women's Ascent up the Organizational Ladder." *Journal of Social Issues* 57:657–74. https://doi.org/10.1111/0022-4537.00234.

Hess, Robert D., and Irene T. Miura. 1985. "Gender Differences in Enrollment in Computer Camps and Classes." *Sex Roles* 13:193–203. https://doi.org/10.1007/BF00287910.

Hill, Catherine, Christianne Corbett, and Andresse St. Rose. 2010. *Why so Few? Women in Science, Technology, Engineering, and Mathematics*. Washington, DC: AAUW.

Hirsh, C. Elizabeth. 2009. "The Strength of Weak Enforcement: The Impact of Discrimination Charges, Legal Environments, and Organizational Conditions on Workplace Segregation." *American Sociological Review* 74:245–71. https://doi.org/10.1177/000312240907400205.

Hochschild, Arlie, and Anne Machung. [1989] 2012. *The Second Shift: Working Families and the Revolution at Home*. London, United Kingdom: Penguin Books.

Jacobs, Jerry A., and Kathleen Gerson. 2001. "Overworked Individuals or Overworked Families? Explaining Trends in Work, Leisure, and Family Time." *Work and Occupations* 28:40–63. https://doi.org/10.1177/0730888401028001004.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *The Quarterly Journal of Economics* 107:35–78. https://doi.org/10.2307/2118323.

King, Joe, Malte Reichelt, and Matt L. Huffman. 2017. "Computerization and Wage Inequality between and within German Work Establishments." *Research in Social Stratification and Mobility* 47:67–77. https://doi.org/10.1016/j.rssm.2016.05.002.

Kristal, Tali. 2013. "The Capitalist Machine: Computerization, Workers' Power, and the Decline in Labor's Share within US Industries." *American Sociological Review* 78:361–89. https://doi.org/10.1177/0003122413481351.

Kristal, Tali, and Yinon Cohen. 2015. "What Do Computers Really Do? Computerization, Fading Pay-Setting Institutions and Rising Wage Inequality." *Research in Social Stratification and Mobility* 42:33–47. https://doi.org/10.1016/j.rssm.2015.07.001. http://linkinghub.elsevier.com/retrieve/pii/S0276562415000463.

Krueger, Alan B. 1993. "How Computers Have Changed the Wage Structure: Evidence from Microdata, 1984–1989." *The Quarterly Journal of Economics* 108:33–60. https://doi.org/10.2307/2118494.

Lapan, Richard T., Angela Adams, Sherri Turner, and Jeanne M. Hinkelman. 2000. "Seventh Graders' Vocational Interest and Efficacy Expectation Patterns." *Journal of Career Development* 26:215–29. https://doi.org/10.1177/089484530002600305.

Legewie, Joscha, and Thomas A. DiPrete. 2014. "The High School Environment and the Gender Gap in Science and Engineering." *Sociology of Education* 87:259–80. https://doi.org/10.1177/0038040714547770.

Levanon, Asaf, and David B. Grusky. 2016. "The Persistence of Extreme Gender Segregation in the Twenty-first Century." *American Journal of Sociology* 122:573–619. https://doi.org/10.1086/688628.

Levy, Frank, and Richard J. Murnane. 1992. "US Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations." *Journal of Economic Literature* 30:1333–81.

Liu, Jonathan, and Doirean Wilson. 2001. "Developing Women in a Digital World." *Women in Management Review* 16:405–16. https://doi.org/10.1108/09649420110411701.

Liu, Yujia, and David B. Grusky. 2013. "The Payoff to Skill in the Third Industrial Revolution." *American Journal of Sociology* 118:1330–74. https://doi.org/10.1086/669498.

Mann, Allison, and Thomas A. DiPrete. 2016. "The Consequences of the National Math and Science Performance Environment for Gender Differences in STEM Aspiration." *Sociological Science* 3:568–603. https://doi.org/10.15195/v3.a25.

Mason, Karen Oppenheim, and Yu-Hsia Lu. 1988. "Attitudes Toward Women's Familial Roles: Changes in the United States, 1977–1985." *Gender & Society* 2:39–57. https://doi.org/10.1177/089124388002001004.

Meyer, Peter B., and Anastasiya M. Osborne. 2005. "Proposed Category System for 1960–2000 Census Occupations." Bureau of Labor Statistics: Technical Report Working Paper 383. Washington, DC: U.S. Bureau of Labor Statistics.

Moss-Racusin, Corinne A., John F. Dovidio, Victoria L. Brescoll, Mark J. Graham, and Jo Handelsman. 2012. "Science Faculty's Subtle Gender Biases Favor Male Students." *Proceedings of the National Academy of Sciences* 109:16474–9. https://doi.org/10.1073/pnas.1211286109.

Mouw, Ted, and Arne L. Kalleberg. 2010. "Occupations and the Structure of Wage Inequality in the United States, 1980s to 2000s." *American Sociological Review* 75:402–31. https://doi.org/10.1177/0003122410363564.

Perlow, Leslie A. 1999. "The Time Famine: Toward a Sociology of Work Time." *Administrative Science Quarterly* 44:57–81. https://doi.org/10.2307/2667031.

Petersen, Trond, and Ishak Saporta. 2004. "The Opportunity Structure for Discrimination." *American Journal of Sociology* 109:852–901. https://doi.org/10.1086/378536.

Phillips, Damon J. 2005. "Organizational Genealogies and the Persistence of Gender Inequality: The Case of Silicon Valley Law Firms." *Administrative Science Quarterly* 50:440–72. https://doi.org/10.2189/asqu.2005.50.3.440.

Reskin, Barbara F. 2000. "The Proximate Causes of Employment Discrimination." *Contemporary Sociology* 29:319–28. https://doi.org/10.2307/2654387.

Reskin, Barbara F., and Debra Branch McBrier. 2000. "Why Not Ascription? Organizations' Employment of Male and Female Managers." *American Sociological Review* 65:210–33. https://doi.org/10.2307/2657438.

Reskin, Barbara F., and Patricia A. Roos. 2009. *Job Queues, Gender Queues: Explaining Women's Inroads into Male Occupations*. Philadelphia, PA: Temple University Press.

Ridgeway, Cecilia L. 1997. "Interaction and the Conservation of Gender Inequality: Considering Employment." *American Sociological Review* 62:218–35. https://doi.org/10.2307/2657301.

Ridgeway, Cecilia L., and Shelley J. Correll. 2004. "Unpacking the Gender System: A Theoretical Perspective on Gender Beliefs and Social Relations." *Gender & Society* 18:510–31. https://doi.org/10.1177/0891243204265269.

Rivera, Lauren A. 2017. "When Two Bodies Are (Not) a Problem: Gender and Relationship Status Discrimination in Academic Hiring." *American Sociological Review* 82:1111–38. https://doi.org/10.1177/0003122417739294.

Rodgers, Willard L., Charles Brown, and Greg J. Duncan. 1993. "Errors in Survey Reports of Earnings, Hours Worked, and Hourly Wages." *Journal of the American Statistical Association* 88:1208–18. https://doi.org/10.1080/01621459.1993.10476400.

Rosenbloom, Joshua L., Ronald A. Ash, Brandon Dupont, and LeAnne Coder. 2008. "Why Are There so Few Women in Information Technology? Assessing the Role of Personality in Career Choices." *Journal of Economic Psychology* 29:543–54. https://doi.org/10.1016/j.joep.2007.09.005.

Salzman, Harold, and Stephen R. Rosenthal. 1994. *Software by Design: Shaping Technology and the Workplace*. New York, NY: Oxford University Press.

Tai, An-Ju R., and Randi L. Sims. 2005. "The Perception of the Glass Ceiling in High Technology Companies." *Journal of Leadership & Organizational Studies* 12:16–23. https://doi.org/10.1177/107179190501200103.

Thurow, Lester C. 1975. *Generating Inequality*. New York, NY: Basic Books.

Tomaskovic-Devey, Donald. 1993. *Gender and Racial Inequality at Work: The Sources and Consequences of Job Segregation. Number 27*. Ithaca, NY: Cornell University Press.

Tomaskovic-Devey, Don, and Sheryl Skaggs. 2002. "Sex Segregation, Labor Process Organization, and Gender Earnings Inequality." *American Journal of Sociology* 108:102–28. https://doi.org/10.1086/344214.

Truman, Gregory E., and Jack J. Baroudi. 1994. "Gender Differences in the Information Systems Managerial Ranks: An Assessment of Potential Discriminatory Practices." *MIS Quarterly* 18:129–42. https://doi.org/10.2307/249761.

Weeden, Kim A. 1998. "Revisiting Occupational Sex Segregation in the United States, 1910–1990: Results from a Log-Linear Approach." *Demography* 35:475–87. https://doi.org/10.2307/3004015. https://doi.org/10.2307/3004015.

Weeden, Kim A. 2002. "Why Do Some Occupations Pay More Than Others? Social Closure and Earnings Inequality in the United States." *American Journal of Sociology* 108:55–101. https://doi.org/10.1086/344121.

Weeden, Kim A., and David B. Grusky. 2005. "The Case for a New Class Map." *American Journal of Sociology* 111:141–212. https://doi.org/10.1086/428815.

Weinberg, Bruce A. 2000. "Computer Use and the Demand for Female Workers." *Industrial and Labor Relations Review* 53:290–308. https://doi.org/10.1177/001979390005300206.

Wilder, Gita, Diane Mackie, and Joel Cooper. 1985. "Gender and Computers: Two Surveys of Computer-Related Attitudes." *Sex Roles* 13:215–28. https://doi.org/10.1007/BF00287912.

Wynn, Alison T., and Shelley J. Correll. 2017. "Gendered Perceptions of Cultural and Skill Alignment in Technology Companies." *Social Sciences* 6:45. https://doi.org/10.3390/socsci6020045.

Xie, Yu, Michael Fang, and Kimberlee Shauman. 2015. "STEM education." *Annual Review of Sociology* 41:331–357. https://doi.org/10.1146/annurev-soc-071312-145659.

Xie, Yu, Alexandra Killewald, and Christopher Near. 2016. "Between- and Within-Occupation Inequality: The Case of High-Status Professions." *The Annals of the American Academy of Political and Social Science* 663:53–79. https://doi.org/10.1177/0002716215596958.

Acknowledgments: Direct all correspondence to Siwei Cheng, assistant professor of sociology at New York University (295 Lafayette St, 4th Floor, New York, NY 10012). The authors acknowledge support from the Summer Research Fund at the Department of Sociology at New York University. The second and third authors contributed equally to the project. We thank Paula England, Kathleen Gerson, Claudia Goldin, Mike Hout, and Yu Xie for helpful comments on earlier versions of this article. This article was presented at the Harvard University Social Demography Seminar. All remaining errors are our own.

Siwei Cheng: Department of Sociology, New York University.
E-mail: siwei.cheng@nyu.edu.

Bhumika Chauhan: Department of Sociology, New York University.
E-mail: bhumikachauhan@nyu.edu.

Swati Chintala: Department of Sociology, New York University.
E-mail: swati.chintala@nyu.edu.