RUNNING HEAD: Class and Earnings Inequality

Social Class and Earnings Inequality

Kim A. Weeden Cornell University

Young-Mi Kim, Cornell University

Matthew Di Carlo, Cornell University

David B. Grusky, Stanford University

November 2005

Keywords: social class, earnings inequality, wage inequality

The research reported here was supported by a National Science Foundation research grant (SBS-9906419) and with discretionary funds from Cornell University and Stanford University. We thank Stephen L. Morgan for his helpful comments and both Stephen L. Morgan and Zun Tang for sharing their implementation of the EG class scheme with us.

The takeoff in earnings inequality is one of the most spectacular social developments in the recent history of the United States. Although an appropriately massive literature on the takeoff has developed among social scientists, it is of course a literature dominated by economists and largely eschewed by sociologists. This state of affairs may seem surprising given that sociologists, rather than economists, have long positioned themselves as uniquely concerned with the (unequal) distribution of valued goods.

Why have sociologists, with few notable exceptions (e.g., Bernhardt et al. 2001; Firebaugh 2003; Morgan & Tang 2005), almost studiously ignored what is arguably one of the most consequential social developments of our time? The main explanation resides, we think, in the discipline's commitment to understanding and measuring inequality with either socioeconomic scales or social class categories, neither of which reference the earnings distribution in any direct or simple way. The unanticipated effect of this measurement decision is to lock sociologists out of what would prove to be one of the key social science literatures in the last quarter-century and thus marginalize the discipline yet further.

Is it possible for sociologists to make a major contribution to the study of earnings inequality even at this late date? There are two approaches that might prove to be fruitful here. First, sociologists might develop uniquely sociological interpretations of the growth in earnings inequality, perhaps by focusing on its political, institutional, or organizational sources (e.g., Morris & Western 1999). This approach is worth pursuing given that some of the long-dominant economic accounts are loosening their hold over the field (see, e.g., Bernstein & Mishel, 1999; DiNardo & Card 2002). Nevertheless, it will be difficult for sociologists to have much impact at this late date, not just because "first-movers" tend to control the development of a literature (i.e., the path dependency effect) but also because one would be hard-pressed, given the unusually long and intensive history of scholarship in the area, to come up with an account that has not been proffered already.

The second approach, and one that we shall take here, is to focus on the consequences of the rise in earnings inequality rather than its sources, thus sidestepping the types of questions that economists conventionally pose. That is, rather than attempting to offer yet another "smoking gun" paper that pits skill-biased technological change against some other, preferred source of the growth in earnings inequality, we will instead explore the implications of rising earnings inequality for the class structure. Although one can perhaps understand why sociologists, interested as they are in social class, might not attend much to ongoing debates on the causes of the takeoff in earnings inequality, one might at least suppose that they would have explored the effects of this major development on the class structure. However, rather little in the way of relevant work on this question has in fact been completed, an omission that we seek to rectify here.

We wish to explore, then, possible changes in the relationship between earnings and social class. This relationship is of interest because class categories are, within at least some traditions of class analysis, conceived as homogeneous "social containers" made up of individuals facing similar constraints and experiencing similar life conditions. Obviously, earnings are very important in defining constraints and life conditions, and one might accordingly ask whether the social classes that sociologists use to study inequality are capturing earnings reliably. In justifying class models and the continuing relevance of class, the association between class and earnings or income has long been emphasized (e.g., Beck 1992; Hout, Brooks, & Manza 1993), although most class analysts also stress that class categories are useful precisely because they signal not just earnings or income but other life conditions as well. If sociologists were interested exclusively in earnings, presumably they would measure it directly rather than pushing it through the fulcrum of class. Nonetheless, there is no disputing that earnings are a fundamental aspect of life conditions, and sociologists should therefore want to know whether our featured measurement tool, the class scheme, is adequately signaling earnings.

We will carry out this analysis in the context of a model that allows for class-like structure to emerge at both the big-class and micro-class levels (see Weeden & Grusky 2005a,b,c; Grusky 2005a,b). The conventional class analyst relies exclusively on a big-class scheme that rests on such categories as professional, manager, clerical worker, sales worker, craft worker, operative, laborer, and farmer. Although we will use conventional big-class categories throughout our analysis, we will also examine how the takeoff played out within the detailed occupations that emerge within each of these categories.

The detailed occupation (e.g., lawyer, secretary, carpenter) is a distinctive social category because it is so deeply institutionalized in contemporary labor markets. This institutionalization takes the form, for example, of (1) workers routinely representing their career aspirations in occupational terms, (2) professional and vocational schools training workers for occupationally-defined skills, (3) professional associations and labor unions forming around occupational designations, and (4) employers constructing, advertising, and remunerating jobs in terms of occupational labels (e.g., Treiman 1977; Wilensky 1966). By contrast, conventional big-class categories are largely statistical constructions fashioned by sociologists and Census Bureau statisticians, and the boundaries around them are therefore analytical in nature rather than deeply institutionalized. The analysis that follows will explore how the takeoff in earnings inequality has affected the homogeneity of such analytical big-class categories as well as that of more deeply institutionalized occupational categories nested within these statistical big classes. For sociologists, these two social forms are the defining structures of modern inequality systems, implying that they should care deeply about how they have developed in response to the takeoff in earnings inequality.

We do not mean to suggest that all prior analyses of the takeoff have invariably ignored big-classes or detailed occupations. To the contrary, many such analyses feature big classes of one sort or another (often referred to as "occupations"), and a handful of analyses bring in detailed

occupations in some way (Groshen 1991; Murphy & Welch 1993). These efforts fall short for our purposes, however, because they analyze either big classes or detailed occupations, but not both. The structure of the takeoff is best understood, we argue, by recognizing that detailed occupations are nested in big classes and then apportioning the earnings inequality into three components:

(1) The between-class component (BC) refers to that portion of inequality in earnings that occurs between big classes.

(2) The between-occupation/within-class component (BO/WC) refers to that portion of earnings inequality that occurs between the detailed occupations that comprise each big class.

(3) The within-occupation component (WO) refers to that portion of earnings inequality that occurs within detailed occupations.

The within-class component estimated in conventional big-class analyses combines BO/WC and WO, and the between-occupation component estimated in conventional occupational analyses combines BC and BO/WC. We will argue that these components are driven by very different mechanisms and are therefore usefully distinguished in understanding trends.

The further virtue of distinguishing these three components is that it allows us to identify whether trends in inequality are playing out in ways that strengthen either big classes or occupations, strengthen neither, or strengthen both. We have graphed in Figure 1 a constellation of changes that, taken together, are distinctly "pro-class" in their implications. In this simple figure, we have depicted a hypothetical labor market with just three big classes (i.e., BC1, BC2, and BC3), and we have further assumed that three occupations (i.e., MC1-MC3, MC4-MC6, MC7-MC9) are nested within each of these big classes. At the individual level, a mere 5 incumbents are allowed to populate each occupation, thus yielding a labor market with 45 workers. The two vertical axes, each of which denote log(wages), reveal a pattern of change in which WO is shrinking, BO/WC is shrinking, and

BC is growing. The second axis (pertaining to "time 2") thus suggests a three-class society marked by substantial between-class differences in log(wages) and trivial within-class differences. The simple point that we seek to make with this figure is that pro-class change involves a decrease in WO, a decrease in BO/WC, and an increase in BC.

We have represented the transition to an occupational (or "micro-class") regime in Figure 2. Although WO is again decreasing in this figure, the other two components of earnings inequality are moving in directions opposite to that shown in Figure 1. That is, between-class differences are now eroding, while between-occupation gaps are now growing. This change effectively yields a labor market with nine small classes rather than three big ones. The occupation principle is accordingly strengthened by a decrease in WO, an increase in BO/WC, and a decrease in BC.

The obvious question that at this point arises is whether a simple increase in the overall inequality (i.e., commensurate increases in WO, BO/WC, and BC) should be treated as at all consequential for either the big-class or micro-class structure. This type of change simply "stretches out" the overall distribution and thus leaves the relative sizes of WO, BO/WC, and BC unchanged (see Allison 1978 for a useful discussion of scale invariance; also, Sen 1973). If an absolutist tack is taken, all that matters is that such change increases the earnings heterogeneity within classes and occupations and accordingly undermines the class principle.[1] The relativist counterargument is that workers in high-inequality societies become inured to inequality, that the social meaning of inequality is accordingly recalibrated to adjust to high dispersion in earnings, and that higher earnings heterogeneity within occupations and classes should therefore be regarded as inconsequential artifacts of living in a high-inequality society (see Blau & Kahn 2002 for a relevant discussion). This counterargument suggests that one should care principally about the relative sizes of WO, BO/WC, and BC. In our own analyses, we shall present evidence on both relative and absolute change, thereby allowing us to remain agnostic and cater to both interpretations.

Sources of Trends

We turn now to a discussion of the mechanisms that underlie changes in each of these components of inequality and that accordingly strengthen or weaken the class principle. As shall be evident, these mechanisms differ across the three components, thus opening up the possibility that rates of change will also differ.

The BC trend

The between-class component, with which we begin our discussion, is well studied within the literature and may therefore be dispensed with quickly. In understanding the sources of trend in this component, it is obviously relevant that entry into different big classes is linked to different levels of schooling, with the professional and managerial classes typically requiring at least a college degree (and typically a more advanced degree as well), the sales and clerical classes typically requiring at least some college education, and the manual classes (i.e., craft, operative, laborer) serving as the default outcome for those with high school degrees or less. We have elsewhere stressed that the mapping between big classes and levels of education is quite imperfect because a great many distinctions embodied in class schemes are only ambiguously and imperfectly related to educational qualifications (e.g., Weeden & Grusky 2005a; Grusky & Sørensen 1998). It is hardly the case, for example, that entry into the craft class invariably requires a high-school degree or that entry into the operative or laboring classes is reserved for those who lack such degrees. Even so, there is clearly a strong association between class and educational qualifications, and BC should accordingly fluctuate with changes in the returns to different levels of schooling. The BC trend line should, then, be quite flat in the 1970s by virtue of the relatively large supply of college-educated workers and then take off in the 1980s as the demand for college-educated workers begins to outstrip the supply (see Katz & Murphy 1992; Levy & Murname 1992).

The foregoing story has classes as epiphenomenal, statistical entities that merely stand in as markers of educational qualifications. Could one imagine a story about BC in which classes are more than such indirect "stand-ins?" This type of story indeed becomes available in societies, such as Sweden, where big classes are deeply institutionalized and their representatives participate formally in collective bargaining. In such an institutional context, big classes become meaningful wage-setting actors, not just statistical categories that indirectly signal changes in market-driven returns to schooling. The BC component would increase, for example, in Sweden insofar as upper-class representatives were especially successful in collective bargaining or lower-class representatives were especially unsuccessful. To some extent, big unions in the U.S. engage in collective bargaining and wage-setting efforts of just this sort, but they do not map onto entire big classes nearly as neatly as in the Swedish case, nor do they command anything approaching the wage-setting power of their Swedish counterparts.

The BO/WC trend

There is no empirical literature upon which we can as readily draw in specifying the likely BO/WC trend (cf. Weeden 2002). We can, however, lay out two simple hypotheses about the BO/WC trend, the first directly related to demand shifts that favor high-skill labor, and the second directly related to the spread of occupationalization and occupational closure. The first story is simply a within-class variant of a standard demand-side account of the takeoff: Namely, we might understand each big class as comprising high-skill and low-skill occupations, and we might further anticipate that the demand for labor within the high-skill occupations began in the 1980s to outstrip the supply of labor and thereby drive up wages. For example, the relative demand for high-skill computer programmers and related technical occupations within the professional class clearly took off during the 1980s, whereas the demand for other less skilled occupations within the professional class (e.g., photographers) appears to have increased more slowly. The increasing demand for

computer programmers and related technical occupations should of course be understood as part of the shift toward an "information economy" driven by technological change and outsourcing.

The distinctive feature of such occupation-based demand shifts is that they can yield long-term wage premia for the expanding occupations. Because high-skill occupations often control entry (via credentials, training requirements, and certification), the supply of labor does not necessarily respond well to the rise in wages, and such demand-induced wage premia can therefore be preserved over the longer run. Indeed, whereas the supply of educational qualifications that govern entry into big classes (e.g., high-school diplomas, college diplomas) is not directly controlled by big classes themselves, the supply of credentials that govern entry into specific occupations (e.g., state bar exams) is to a greater extent controlled by the affected professions. This control can be parlayed into slow-growth policies that advantage current incumbents by keeping wages high. There is no mechanism, by contrast, through which the "professional class" can similarly restrict the number of bachelor's degrees awarded.

The occupation-based variant of the demand-shift story can therefore be pitched in a way that references the capacity of some occupations to control the supply of labor. We can add further dynamism to such an account by recognizing that some occupations are progressively developing this capacity to control the market. For many occupations, a key objective of the affiliated occupational or professional association is not just to secure control over the certification of labor (e.g., the control of the American Medical Association over the training of doctors), but also to gain monopoly control over the provision of its core product or service (e.g., the control of lawyers over adjudicating legal disputes). This process of securing closure can raise the wages of occupational incumbents by (1) ceding control over the supply of labor to the occupation, and (2) reducing competition from other potential providers of the product or service and therefore increasing demand (see Weeden 2002; Abbott 1988). If closure movements were a random process to which all

occupations, both high-skill and low-skill alike, were equally subjected, then they wouldn't have any implications for within-class inequality. However, high-skill occupations appear to have been especially successful in securing closure of this sort, thereby raising their already high wages yet higher. This process has been especially apparent within the professions (i.e., "professionalization") and other nonmanual big classes.

If professionalization has historically been the method of choice for securing the rewards of social closure within the nonmanual sector, craft unionization is the corresponding method of choice within the manual sector. The latter method has fared less well of late than has professionalization. As craft unions weaken, it has become increasingly difficult for them to maintain control over the supply of labor and to capture the rent that such control makes possible, and the wage gap between the historically privileged craft occupations and all others has accordingly begun to close. The BO/WC trend may therefore be suppressed by such countervailing "deoccupationalizing" effects in the manual big classes.

The WO trend

The final component of interest is the within-occupation variance (WO). In the economics literature, the closest analogue to the within-occupation variance is what has conventionally been termed the "within-group" variance, where this refers to the residual variance in earnings that falls out of standard Mincer-type human capital regressions (using such variables as age, years of schooling, and interactions between age and years of schooling). Although the groups that are referenced by such methods are statistical cross-classifications rather than institutionalized occupations (cf. Groshen 1991), it is still worth examining the ways in which the within-group variance is presumed within this literature to be generated.

The stylized fact with which this literature begins is that most of the growth in overall wage inequality is attributable to growth in such residual inequality (e.g., Acemoglu 2002; DiNardo, Fortin,

& Lemieux 1996; Juhn, Murphy, & Pierce 1993; Katz & Autor 1999; cf. Card & DiNardo 2003; Lemieux 2003). This growth is frequently attributed, in turn, to the rising demand for skill, where the skills of interest are not indexed by observable schooling but by unobservable social, cognitive, or personality traits. Although the presumption is that the demand for skilled labor started expanding in the 1970s, the effect of this expansion on the wage premium for observable schooling was suppressed until the 1980s, when finally the supply of college-schooled labor no longer kept pace with the growing demand for it. By contrast, the relative supply of unobservable skills may be presumed to be constant over time, meaning that the wage premium for such skills should steadily expand and generate a long-term upward trend in residual inequality over the 1970s and beyond (see especially Acemoglu 2002).

The rising demand for unobservable skills, which we treat as exogenous here, can itself be attributed to any number of causes, including SBTC and shifts in the structure of final demand (also see Stone 2004; Frank & Cook 1995). Whatever its source, a rising payoff to unobservable skills will generate a growth in within-group inequality, just as a rising payoff to observable skills generates a growth in between-group inequality. This type of account is of course attractive because it is consistent with competitive wage theory.

We are not averse to porting over this economic account and using it to explain trends in within-occupation inequality (WO). It is possible, however, that much of what appears to be growing demand for high skill labor within the residual statistical groups of Mincer-type regression models is just growing demand for particular high-skill occupations, not growing demand for high-skill tasks within all occupations. If so, we need only invoke again a story about increasing demand for high-skill occupations, a story already outlined in the preceding section on the BO/WC trend. That is, once the BO/WC inequality is stripped out of the residual, it may well be that the remaining WO component suggests much less in the way of growing demand for high-skill labor.

We of course appreciate that much residual inequality is likely to remain within detailed occupations. It is entirely plausible, moreover, that increases in such inequality are partly attributable to rising demand for unobservable skill, just as the standard economic account would have it. At the same time, this account would appear to be most appropriate in explaining a takeoff in the nonmanual sector, where the introduction of computers and related technical innovations has transformed the workplace in especially revolutionary (and skill-demanding) ways. We will therefore take care in the analyses that follow to examine the evolution of WO in the manual and nonmanual sectors separately.

Caveats

We will not attempt in this paper to examine whether these, or any other, mechanisms are in fact playing out as just described. Rather, our purpose in rehearsing the mechanisms is simply to suggest that quite different processes may generate change in BC, BO/WC, and WO, not just change in the absolute size of these components but also change in their relative size. We focus in this paper on merely describing these trends, which remain uncharted even at this late date, and then assessing whether they are pro-class or pro-occupation in their implications.

We begin this task by exploring overall trends for each of these three components. However, we turn thereafter to exploring trends in BO/WC and WO *within* each of the big classes, an approach that is useful because the overall trends potentially differ by class. The preceding discussion implies, in particular, that upward trends in BO/WC and WO may be more pronounced in the nonmanual than manual classes.

Data and Methods

The following analysis relies on earnings data collected in the May supplement of the Current Population Surveys [CPS] from 1973 to 1978 and in the "outgoing rotation group" (ORG) supplements of the monthly CPS from 1979 to 2004 (Bureau of Labor Statistics [BLS], years vary).

These surveys both contain information on the usual pay of the main job that the individual held in the week prior to the survey. In constructing our samples from the May/ORG CPS files, we include all wage and salary workers between the ages of 16 and 65 who report positive hours worked at their main jobs and for whom a valid occupation code is available.

Data processing

To prepare the May/ORG files for the analysis, we were required to make a host of technical and conceptual decisions that affect the measurement of inequality and, potentially, observed trends in inequality. We of course sought to apply current "best practice" when making these decisions. In the absence of such a consensus, we carried out analyses using two or more practices and compared the results, usually with the outcome that the main conclusions were unaffected. We will therefore report but one set of results in the body of the paper and relegate the tedious task of reporting results of alternative specifications to footnotes or the Appendix.

The first set of relevant technical decisions pertains to the measurement of earnings. In the May/ORG CPS, workers who indicate that they are paid on an hourly basis are asked to report their usual hourly wages, and workers who indicate some other pay periodicity report their usual weekly wages. We follow here conventional practice of calculating hourly wages for non-hourly workers by dividing weekly wages by the number of hours usually worked at the main job. We then convert hourly wages into constant 2000 dollars using the personal consumption expenditures index from the National Income and Products Accounts (Bureau of Economic Analysis, 2005).[2]

As is also conventional practice, we exclude workers whose hourly wages fall below $1 or above $100 in 1979 dollars (e.g., Angrist & Krueger 2000; Card & Dinardo 2002; Mishel, Bernstein, & Schmitt 2001; cf. Bollinger & Chandra 2005). We also perforce exclude self-employed workers, given that these workers are not asked the earnings questions in the ORG supplements.[3] After these exclusions, our samples contain on average 20,254 men and 15,776 women per year for the May

surveys, and 64,048 men and 62,181 women per year for the ORG surveys, thus yielding a total of 1,618,174 men and 1,587,005 women.

The second data processing decision of potential consequence pertains to the use of edited or unedited earnings. The BLS allocates earnings to the growing percentage of workers who, for whatever reason, chose not to answer the earnings questions. The complication that arises in using such edited earnings is twofold: first, they are not available in the May 1973-1978 surveys; and, second, the BLS changed in the mid-1990s the hot-deck procedure it used to allocate earnings (see Hirsch & Schumacher 2003; Lemieux 2003). Because both of these procedural shifts could affect observed trends in wage inequality, we follow Lemieux (2003) and others in relying on unedited earnings. This forces us to exclude the 1994-95 ORG files in which workers with edited earnings are impossible to identify (see Lemieux 2003, pp. 6-7).

The third decision of potential consequence centers around the effects of topcoding censored wages of high-earning workers. In the ORG surveys, topcoding does not have much impact on hourly workers, because very few have hourly wages greater than the topcode value (i.e., $99.99). However, topcoding does affect between 1% and 4% of workers who are paid on a nonhourly basis in most years, and it affects an even higher percentage of such nonhourly workers in the 1980-1983 surveys. The topcode value also changes over the sample period. For example, weekly earnings over $999 (in nominal dollars) are censored in the 1973-1980 files, whereas weekly earnings over $2,884 are censored in the 1998-2004 files. We follow conventional practice here in adjusting for topcoding by multiplying topcoded earnings by 1.4 (see, e.g., Lemieux 2003; Card and DiNardo 2002).[4]

The fourth and final data-processing decision involves the treatment of part-time workers. Although such workers are typically excluded from analysis, a small but growing contingent of labor economists advocates weighting workers by the number of hours they usually work, thereby

obtaining a wage distribution representative of all hours worked in the economy (e.g., DiNardo, Fortin, & Lemieux 1996; Lemieux 2003; Card & DiNardo 2002). The latter practice, which we follow here, may be understood as a good compromise between excluding part-time workers and simply ignoring the distinction between part-time and full-time work.[5]

Occupation and Class Schemes

The measurement of class and occupation raises a fresh set of complications. Because of the great many competing big-class and occupation schemes, good practice requires using multiple such schemes, an approach we follow here. We describe below the two occupation schemes and two big-class schemes that we have applied.

For most of our reported analyses, we have coded occupations in terms of the Standard Occupation Classification (SOC), thus allowing the classification scheme to change as the BLS introduced new classification systems in the 1983, 1992, and 2003 surveys. We also examine trends with a uniform scheme created by back-coding all CPS data between 1983 and 2004 into the 1970 SOC codes (see Weeden 2005b, c). This backcoded 1970 scheme is created in three steps. First, we translate the 2000-basis data into the 1990 scheme by (1) multiplying each 2000-basis record by the number of 1990-SOC codes that contribute to the 2000 code (U.S. Bureau of the Census 2005), and (2) assigning sex-specific weights to each record in the resulting 1990-basis expanded data set.[6] This weight equals the proportion of the 2000 code that is drawn from the constituent 1990 code. We calculate these weights by matching the CPS Basic files in 2000-2002 to the CPS Extract files that contain the 2000-basis occupation code for the same individuals. Second, we translate all 1990-basis data (including, now, the 2003 and 2004 CPS) into the 1980 scheme, a translation that involves a quite trivial set of recodes. Third, we translate all 1980-basis data (now including data from 1992-2004) into the 1970 code, using a weighted back-coding procedure analogous to that described above for the 2000-to-1990 recodes (see also Weeden & Grusky 2005a,b,c).

Which occupation scheme is to be preferred? The back-coding procedure used to construct the constant, 1970-basis scheme introduces measurement error, which in turn leads to an inflated estimate of the amount of within-occupation inequality in a given year. The indigenous occupation scheme, by contrast, represents the best available characterization of the occupational structure in a given year, thus presumably allowing for the most accurate estimate of absolute levels of inequality in each year. However, where the goal is to obtain estimates of *trend* in the various components of inequality, one ought not necessarily default to accuracy of this sort. Indeed, if each new occupational classification scheme comes on the heels of real change in institutionalized boundaries of the occupational structure it purportedly represents, analysis of the data arrayed in the indigenous scheme will offer a better assessment of trends in micro-class inequality. The 1970-basis scheme, by contrast, will overestimate the increase in within-occupation inequality and, conversely, underestimate the increase in between-occupation inequality, at least if we assume that the earnings profiles of newly formed occupations diverge from those of their "parent" occupation. If, on the other hand, each new occupation scheme merely improves the measurement of an underlying occupational structure that has remained largely constant, the indigenous occupation scheme will overestimate the increase in between-occupation inequality. It is entirely plausible that both accounts are partly true: That is, some of the changes in a new occupation scheme may capture newly institutionalized occupational boundaries (e.g., systems analysts), whereas others may capture occupational distinctions that existed in the past but had simply been ignored by the classifiers (e.g., distinctions among types of managers).

Because there is no solution to this dilemma, we have again estimated trends in the various components of inequality with both indigenous and 1970-basis schemes, again allowing us to gauge the sensitivity of our results to our measurement decisions.[7] In the interest of brevity, we report only the indigenous-scheme decompositions in the body of the paper, noting any cross-scheme

differences in footnotes. We also present an illustrative decomposition of inequality using the 1970-basis codes in the Appendix.

We evaluate these two occupation schemes against two conventional big-class representations of the site of production: the Erikson-Goldthorpe (EG) scheme, and the Featherman-Hauser (FH) scheme. The EG scheme has become the *de facto* standard in European analyses of inequality, and it also has its share of users, if not advocates, in North American analyses (e.g., Morgan & Tang 2005; Manza & Brooks 1999). We use the commonly applied 7-class variant of EG, which can be implemented, albeit imperfectly, with occupation codes (Erikson & Goldthorpe, 1992, pp. 35-47). However, two of the seven classes, the petty bourgeoisie and farmers, are perforce excluded from our analyses, because they contain only self-employed workers. The five remaining EG categories are service workers, routine nonmanuals, skilled craft workers, unskilled manual workers, and agricultural workers.

The task of implementing the EG scheme was complicated by the absence of standard algorithms for translating SOC codes into EG categories. For this paper, we relied heavily on Morgan and Tang's (2005) translation of 1980 and 1990 SOC codes into EG classes.[8] Using the Morgan-Tang algorithm and the double-coded 2000-2002 CPS Basic files (see above), we created our own algorithm for translating 2000 SOC codes into EG categories by (1) identifying all of the 1990 SOC occupation codes that contributed to a given 2000 SOC code, and (2) calculating the modal EG class for each SOC 2000 code based on the relative size of each 1990 contributor and its EG class assignments. We used an analogous algorithm to translate 1970 SOC codes into the EG scheme.[9] We then created one set of EG codes for the 1973-1982 CPS files, using the 1970-basis algorithm to assign an EG category to each record, and two sets of EG codes for the 1983-2004 CPS files, one using the SOC-specific algorithm to translate the indigenous occupation codes into EG codes, and one using the 1970-basis algorithm to translate the constant, 1970-basis occupation

codes into EG codes.

For all its popularity, the EG scheme has not yet supplanted alternatives based on aggregate Census categories, of which arguably the most prominent is the scheme popularized by Featherman and Hauser (1978). Indeed, our prior research reveals that the FH scheme outperforms the EG scheme in capturing class-based heterogeneity in a multitude of outcome variables (including, but not limited to, income), a finding that we attributed to the greater detail of the FH scheme as well as its history as a scheme devised specifically for the idiosyncrasies of the U.S. case and U.S. occupational classification systems (Weeden & Grusky 2005a,b,c; Weeden 2005). Consistent with this prior work, we feature the FH scheme in the present paper, assessing trends in wage inequality between and within the categories of this scheme. We discuss the comparison of FH and EG results in the Appendix.

The body of the paper will, however, rest on results for the FH scheme. We begin with the standard, 12-class version of FH and excise the two self-employed classes (i.e., self-employed professionals, self-employed managers). The remaining 10 FH categories are employed professionals, employed managers, sales workers, clerical workers, craft workers, operatives, service workers, laborers, farmers, and farm laborers.

Methods

We rely on two measures of inequality, the variance of (log)wages and Theil's index, both of which can readily be decomposed into BC, BO/WC, and WO components. For some analyses, we also calculate the log of the 90/10 ratio, which is less sensitive to topcoding than our other measures. We of course provide separate decompositions of inequality for men and women throughout.

The total variance in log(wages) can be decomposed by fitting two models: the first regresses log(wages) on the categories of a big-class scheme, and the second regresses log(wages) on the

categories of a detailed occupation scheme, where occupations are nested in big classes. These two models allow us to implement our desired three-way decomposition: The BC component is the total variance in log(wages) minus the variance of the residuals of the big-class model; the BO/WC component is the difference between the variance of the residuals of the occupation and big-class models; and the WO component is the variance of the residuals from the occupation model. These definitions are easily formalized:

$$Var_{BC} = \frac{\sum_{i=1}^{N}(y_i - \bar{y})^2}{N-1} - \frac{\sum_{i=1}^{N}\left((\hat{y}_i|c_i - y_i) - \overline{(\hat{y}_i|c_i - y_i)}\right)^2}{N-1},$$

$$Var_{WC/BO} = \frac{\sum_{i=1}^{N}\left((\hat{y}_i|o_i - y_i) - \overline{(\hat{y}_i|o_i - y_i)}\right)^2}{N-1} - \frac{\sum_{i=1}^{N}\left((\hat{y}_i|c_i - y_i) - \overline{(\hat{y}_i|c_i - y_i)}\right)^2}{N-1}, \text{ and}$$

$$Var_{WO} = \frac{\sum_{i=1}^{N}\left((\hat{y}_i|o_i - y_i) - \overline{(\hat{y}_i|o_i - y_i)}\right)^2}{N-1},$$

where c_i is a vector of dummy variables indexing the big class of individual i, o_i is a vector of dummy variables indexing the occupation of individual i, and N indicates the sample size.

The decomposition of Theil's index is frequently decomposed into two components, but less frequently decomposed into three components, as required here. This three-stage decomposition takes the following form (see, e.g., Akita 2001):

$$T_{BC} = \sum_{c=1}^{C}\frac{Y_c}{Y}\ln\left(\frac{Y_c}{Y}\Big/\frac{N_c}{N}\right),$$

$$T_{WC/BO} = \sum_{c=1}^{C}\left(\frac{Y_c}{Y}\right)\sum_{o=1}^{O}\left(\frac{Y_{co}}{Y_c}\right)\ln\left(\frac{Y_{co}}{Y_c}\Big/\frac{N_{co}}{N}\right),$$

$$T_{WO} = \sum_{c=1}^{C}\sum_{o=1}^{O}\left(\frac{Y_{co}}{Y}\right)\sum_{i=1}\left(\frac{Y_{coi}}{Y_{co}}\right)\ln\left(\frac{y_{coi}}{Y_{co}}\Big/\frac{1}{N_{co}}\right),$$

where y_{coi} is the hourly wage of individual *I,* N_{co} is the sample size of occupation *o* in class *c*, N_c is the sample size of class *c*, and $Y_{co,}$ $Y_{c,}$ and Y are specified as follows:

$$Y_{co} = \sum_i y_{coi} \text{ ; } Y_c = \sum_o \sum_i y_{coi} \text{ ; and } Y = \sum_c \sum_o \sum_i y_{coi} .$$

The total Theil, *T,* is merely the sum of the three components:

$$T = T_{BC} + T_{WC/BO} + T_{WO} .$$

We will also examine the BO/WC and WO components within each big-class separately. In implementing these decompositions, we again rely on the variance of log(wages) and Theil's index, but here our task is simplified because there are but two components of inequality, BO/WC and WO. The variance within each big class can be decomposed by regressing logged wages on the occupations (measured with O-1 dummy variables) that constitute each big class. The variance of the residuals from this model indicates the amount of within-occupation inequality in that class (and in that year), while the difference between the total variance and the residual variance represents the amount of between-occupation inequality. The decomposition of the Theil index merely requires stratifying the data by class and then calculating the two components in the usual way.

Although all analyses were carried out with both variance-based and Theil decompositions, we will only present a full set of results for the overall trend in inequality, thereafter focusing on the variance of (log)wages. For readers who wish to make more comprehensive comparisons, the decompositions of the Theil index are presented in Figures A1 and A2, and a brief discussion of those results is provided in the Appendix as well.

Results

We set the stage for our decomposition by presenting trends in overall wage inequality from 1973 to 2004 (see Figure 3). These results are calculated with the variance of log(wages), the Theil index, and the log of the 90/10 ratio, the latter normalized as usual by dividing by 2.56 (see Card &

DiNardo 2002, p. 746).[10] We observe for all measures the now-famous pattern of stability in the 1970s, takeoff in the 1980s, and somewhat slower accretion throughout the 1990s and beyond.

The first set of decompositions are presented in Figure 4 (for men) and Figure 5 (for women). As may be recalled, these results are based on a three-way decomposition of the variance of log(wages), using indigenous SOC occupation schemes and FH classes (see Appendix for a discussion of supplementary results). Because the relative rate of change can be difficult to gauge via inspection alone, we report in Table 1 the estimates secured by imposing, for men and women separately, a linear trend on the variance in logged wages. The results under the heading "Model 1" pertain to coefficients from models that fit a single slope and intercept to the trend lines, while the results under the heading "Model 2" pertain to coefficients from models that again fit a single slope but now allow the intercept to shift with each change in the occupational classification scheme (in 1983, 1992, and 2003). The former coefficients correspond to the assumption that new classification schemes are introduced in response to real changes in institutionalized occupational boundaries (i.e., "changing occupational structure"). If such changes are indeed real, then any increase in the explanatory power of these more refined classifications bespeaks a meaningful change in how inequality is apportioned within and between occupations. By contrast, the "Model 2" coefficients correspond to the assumption that new classification schemes capture occupational distinctions that were always present, but were ignored in the more primitive classifications of prior years (i.e., "constant occupational structure"). It is appropriate under this assumption to parse out the classification effects and fit a pooled slope after doing so.

We noted at the outset of our paper that scholars of income inequality might reasonably attend to changes in both the absolute and relative sizes of these three variance components. The results of Table 1 pertain of course to absolute changes alone. We thus present in Table 2 measures of relative change that condition on the models of Table 1. These measures, calculated for 1973 and

2004, are simply the percentage of the total variance in log(wages) that each component represents.[11]

We can at long last turn to answering the seemingly simple questions with which we began. If the BC trendline is considered first, here one has to be struck by the steep takeoff in the 1980s and the steady increase thereafter, especially for men. The between-class component accounts for 45.2% of the total change in male inequality (Model 1, Table 1), and the between-class share of the total male variance increased, moreover, from 17.2% in 1973 to 27.1% in 2004 (Model 1, Table 2). For women, the pooled BC trendline is clearly less steep, and the BC component accordingly accounts for less of the total change in inequality (see Table 1).[12] Although we did not anticipate such a pronounced difference by gender, the general result that the BC component is fast increasing is hardly surprising. If one accepts our interpretation that big classes are effectively education groups in disguise, we are merely reexpressing here the well-known increase in the payoff to schooling.

We consider next the WC/BO trendline. In this case, our results cannot be straightforwardly read off from prior research, as big classes and detailed occupations have not before been analyzed in tandem. We find that, for males and females alike, the WC/BO trendline does not show the characteristic steep takeoff in the early 1980s, nor does it move sharply upward at any point thereafter. The rise is instead a slow and gradual one (except for a classification-induced uptick in 2003) that translates into a slightly increased share of the total variance by 2004 (see Table 2).[13] The WC/BO component accounts overall for 11-17% of the total growth in male inequality and for 19-20% of the total growth in female inequality (see Table 1). As with the BC component, we thus find that the WC/BO component is increasing both in absolute and relative terms, albeit at a pace rather slower than the BC pace. It is possible, of course, that this reduced overall rate conceals the emergence of substantial inter-occupational differences in the nonmanual sector, where occupational closure movements have been most successful and the demand for high-skill "information economy" occupations most pronounced. We will take on this hypothesis subsequently.

The third series in our analysis, the WO series, likewise fails to match the fast-paced growth of the BC series. Indeed, the male WO series suggests a one-time takeoff in the early 1980s, followed by near stasis thereafter. The female series, by contrast, shows clear evidence of ongoing increase after the early 1980s, with a steeper overall slope coefficient as a result. The relatively slow pace of change, especially for men, means that only a minority share of the total change is attributable to a growth in WO (see Table 1), a result that is inconsistent with the long-standing claim that rising "within-group" inequality is the driving force behind the growth of inequality (cf. Lemieux 2003). This result is consistent with our suggestion that much of what conventional Mincer-type formulations represent as growth in within-group inequality is in fact growth in the WC/BO component.

The relatively slow pace of change in the WO component translates, furthermore, into a decline in the percentage of the total variance that is generated within occupations. It may be recalled in this context that the BC and WC/BO components both increased in absolute as well as relative terms. Clearly, these two components had to be taking relative share from the WO component, and indeed Table 2 reveals just this.[14] We find that, when Model 1 estimates are used, the within-occupation share declines from 68.6% to 57.9% for males and from 61.8% to 55.5% for females. The decline is less dramatic, however, under Model 2 (and indeed it disappears altogether for females under that model).

The next set of analyses allow us to examine whether the between-occupation trends are playing out differently in different big classes. In Figures 6a-6d, we graph the between-occupation variance in each of 8 FH big-classes, presenting the manual and nonmanual trendlines separately. We had argued, it may be recalled, that the upward trend in the between-occupation component is likely to be especially pronounced in the nonmanual classes, not just because such classes contain precisely the high-skill occupations (e.g., computer programmer) that might profit from demand

explosions but also because they are the home ground of occupationalization movements that allow high-skill occupations to control supply and demand forces and garner "rent" as a result.

The results in Figures 6a-6d are partly consistent with such claims. For women and men alike, the nonmanual trendlines tend to move upward at a faster pace than the manual trendlines, the latter being mired for the most part at very low levels over the last 30 years. The most dramatic takeoff in between-occupation variance occurs in the nonmanual big-class of sales workers. When we examine the particular occupations that, within this class, are rising or falling in earnings, we find much evidence of high-skill winners (e.g., sales engineers, stock and bond salesmen, advertising agents) and low-skill losers (e.g., retail sales clerks, cashiers, sales representatives), just as a demand-shift hypothesis would have it.

The main exception to our hypothesis is of course the service class (for men). If we again examine the constituent trendlines, we find that such low-skill occupations as maids, waiters, food service workers, and food counter workers are declining in wages, while such high-skill occupations as firemen, marshals, and sheriffs are increasing in wages. This result seems consistent on the face of it with a demand-shift hypothesis.[15] For the most part, nonmanual occupations encompass much skill differentiation and thus provide the home ground for a demand-shift hypothesis, yet the FH service class happens also to be an amalgam of occupations that is quite heterogeneous in skill requirements. It again bears noting in this regard that some of these high-skill occupations have gained (partial) control over the certification process and can therefore resist any massive wage-correcting influx of new workers.

We have presented the corresponding trendlines for the within-occupation components in Figures 7a-7d. As may be recalled, we had argued here that rising demand for unobservable skills should play out principally in the nonmanual sector, where the introduction of computers and related technical innovations have transformed the workplace in especially revolutionary (and skill-

demanding) ways. By contrast, one would be hard-pressed to cite technical or other innovations in craft, operative, service, or laboring occupations that would generate new premia to unobservable skills and hence produce increased wage differentiation.

For both men and women, we find that the results are quite consistent with this type of account, indeed almost eerily so. We find, for example, high and increasing WO trendlines in Figures 7a and 7c and low and stable WO trendlines in Figures 7b and 7d. Although our results are consistent, then, with the assumption of rising demand for unobservable skills, we have to confess that our close inspection of the data reveals a great many occupations with exploding variances that cannot, it would seem, be explained in these terms. As one example, the variance for religious workers took off during this period, yet we doubt that the much-discussed rise of mega-churches is widespread enough to be driving this result nor, even if it were, would it necessarily be interpretable as a consequence of new technologies conducive to mega-churches. This example is simply to make obvious that a far more formal analysis would be required to gain any leverage on the mechanisms at work.

Discussion

It may be useful to conclude by asking how our results speak to the simple questions with which we led off this paper. Has the ongoing increase in earnings inequality played out in pro-class ways? Does the takeoff in inequality instead support the rise of micro-classes? Can both types of class principles be simultaneously supported? Or is earnings inequality becoming an increasingly individualistic affair that occurs within groups rather than between them? By posing questions of this sort, we have sought to move away from the field's long-standing focus (obsession?) with the *causes* of earnings inequality, moving instead to questions about its *effects* on the structure of social classes. It might be supposed that sociologists would have by now explored the effects of the

historic takeoff in earnings inequality on the class structure. Surprisingly, they have not yet done so, save in the most tentative ways (e.g., Kim & Sakamoto 2005).

We thus begin by asking whether the results reported here might be regarded as supportive of big-class formulations (see Figure 1). Although there is, as we see it, much here that big-class enthusiasts can embrace, our data are not working exclusively in support of the big-class principle. The core results to which big-class adherents will inevitably point are the growing absolute size of the between-class component and the growing between-class share of the total variance. Important though these results are, it must also be noted that inter-occupational cleavages within some big classes (e.g., service) are growing as well, a result that is surely inconsistent with simple big-class stories. The double-edged sword of rising earnings inequality is *simultaneous* growth in distinctions between classes and within them. These within-class cleavages have become so prominent in some classes, such as the service class, that it has to be asked whether the presumed class is any longer a single class, if ever it was.

How, then, does the micro-class story fare? Here again, one can find results that are consistent with the micro-class story, and not just the steady growth in within-class occupational distinctions (as noted above). It is additionally relevant that the within-occupation component to inequality is declining in relative size. Obversely, we can conclude that the class principle, as expressed via big classes and micro classes together, is accounting for a growing share of the total variance in earnings. At least as regards income data, a new type of dual closure appears to be slowly emerging, a form of inequality in which social groups at two distinct levels figure increasingly prominently.

This line of argumentation ought not be pushed too hard. It bears noting that, even after 25 years of growth in the "class principle," a majority of the total variance in earnings still occurs within rather than between classes. It should also be recalled that the decline in the within-occupation

variance is playing out almost exclusively within the manual sector. The ever-growing size of the nonmanual sector, where the within-occupation variance is instead increasing, may place real limits on the long-term growth of the class principle.

References

Abbott, A. (1988). The system of professions: An essay on the division of expert labor. Chicago: University of Chicago Press.

Acemoglu, D. (2002). Technical change, inequality, and the labor market. *Journal of Economic Literature*, 40, 7-72.

Akita, T. 2001. Regional income inequality in China: a two-stage nested inequality decomposition analysis. Unpublished manuscript. International University of Japan.

Allison, P.D. (1978). Measures of inequality. *American Sociological Review,* 43, 865-80.

Angrist, J. D., & Krueger , A. B. (2000). Empirical strategies in labor economics. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (volume 3A). Amsterdam: Elsevier Science.

Beck, U. (1992). Risk society: Towards a new modernity. M. Ritter (Trans.). London: Sage.

Bernhardt, A., Morris, M., Handcock, M.S., & Scott, M. A. (2001). Divergent paths: Economic mobility in the new American labor market. New York: Russell Sage.

Bernstein, J. & Mishel, L. (2001). Seven reasons for skepticism about the technology story of U.S. wage inequality. In I. Berg & A. L. Kalleberg, *Sourcebook of labor markets: Evolving structures and processes*. New York: Plenum.

Blau, F.D. & Kahn, L.M. (2002). At home and abroad: U.S. labor market performance in comparative perspective. New York: Russell Sage Foundation.

Bollinger, C. R., & Chandra, A. (2005). Iatrogenic specification error: A cautionary tale of data cleaning. *Journal of Labor Economics*, 23, 235-57.

Bureau of Economic Analysis (2005). National Income and Products Accounts, personal consumption expenditures. Retrieved November 23, 2005 from http://www.bea.gov/bea/dn/nipaweb/Index.asp.

Bureau of Labor Statistics. (Years vary.) May and ORG Supplements to the Current Population Survey. Available in machine-readable form from National Bureau of Economic Research.

Card, D., & DiNardo, J. E. (2002). Skill-biased technological change and rising wage inequality: Some problems and puzzles. *Journal of Labor Economics*, 20(4), 733-83.

Davis, J.A., T.W. Smith, and P. Marsden. (2004). *GSS (General social survey) 1972-2002 cumulative data file*. Chicago: NORC.

DiNardo, J., Fortin, N. M. & Lemieux, T. (1996). Labor market institutions and the distribution of wages, 1973-1992: A semi-parametric approach. *Econometrica*, 65, 1001-46.

Erikson, R. E., & Goldthorpe, J. H. (1992). *The constant flux: A study of class mobility in industrial societies.* Oxford: Oxford University Press.

Featherman, D. L., & Hauser, R. M. (1978). *Opportunity and change.* New York: Academic.

Firebaugh, Glenn (2003). *The new geography of global income inequality*. Cambridge: Harvard University Press.

Frank, R.H., & Cook, P. J. (1995). The winner-take-all society. New York: Free Press.

Groshen, E. (1991). Sources of intra-industry wage dispersion: How much do employers matter? *The Quarterly Journal of Economics,* 106, 869-84.

Grusky, D. B. (in collaboration with Gabriela Galescu). 2005a. Foundations of class analysis: A Durkheimian perspective. *Approaches to class analysis*, edited by Erik O. Wright. Cambridge: Cambridge University Press.

Grusky, D. B. 2005b. Is Durkheim a class analyst? *The Cambridge companion to Durkheim*, edited by Jeffrey Alexander and Phil Smith. Cambridge: Cambridge University Press.

Grusky, D.B., & Sørensen, J. B. (1998). Can class analysis be salvaged? *American Journal of Sociology*, 103, 1187-1234.

Hirsch, B., & Schumacher, E. (2004). Match bias in wage gap estimates due to earnings imputation. *Journal of Labor Economics*, 22, 689-722.

Hout, M., Brooks, C., & Manza, J. (1993). Persistence of classes in postindustrial societies. *International Sociology,* 8, 259-78.

Juhn, C., Murphy, K. M. & Pierce, B. (1993). Wage inequality and the rise in returns to skill. *Journal of Political Economy*, 101, 410-42.

Katz, L. F., & Autor, D. H. (1999). Changes in the wage structure and earnings inequality. In O. C. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (volume 3). Amsterdam: Elsevier Science.

Katz, L.F., & Murphy, K. (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107, 35-78.

Kim, C., & Sakamoto, A. (2005). Bringing inequality back in: The role of occupational structure in growing wage dispersion. Paper presented at the RC 28 Meetings of the International Sociological Association.

Lemieux, T. (2003). Residual wage inequality: A re-examination. Unpublished manuscript. University of British Columbia.

Levy, F. & Murname, R. J. (1992). U.S. earnings levels and earnings inequality: A review of recent trends and proposed explanations. *Journal of Economic Literature*, 30, 1333-81.

Manza, J., & Brooks, C. (1999). *Social cleavages and political change: voter alignments and U.S. party coalitions.* Oxford: Oxford University Press.

Mishel, L., Bernstein, J. & Schmitt, J. (2001). *State of working America: 2000/2001.* Ithaca, NY: Cornell University Press.

Morgan, S. L., & Tang, Z. (2005). Social class and workers' rent, 1983-2001. Unpublished manuscript. Department of Sociology, Cornell University.

Morris, M., & Western, B. (1999). Inequality in earnings at the close of the twentieth century. *Annual Review of Sociology*, 25, 623-57.

Murphy, K., & Welch, F. (1993). Occupational change and the demand for skill. *American Economic Review*, 83, 122-26.

Schmitt, J. (2003). Creating a consistent hourly wage series from the Current Population Survey's Outgoing Rotation Group, 1979-2002. Unpublished manuscript. Center for Economic and Policy Research.

Sen, A. (1973). On economic inequality. Oxford: Oxford University Press.

Stone, K. (2004). From widgets to digits: Employment regulation for the changing workplace. Cambridge, UK: Cambridge University Press.

Treiman, D.J. (1977). Occupational prestige in comparative perspective. New York: Academic Press.

Weeden, K. 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.

Weeden, K. A. (2005a). Big-class politics or occupation politics? Unpublished manuscript. Department of Sociology, Cornell University.

Weeden, K. A. (2005b). Algorithm for translating 1980 and 1990 U.S. census occupation codes into 1970 codes. Department of Sociology, Cornell University.

Weeden, K. A. (2005c). Algorithm for backcoding 2000 census occupation codes into 1990 census occupation codes. Department of Sociology, Cornell University.

Weeden, K. A., & Grusky, D. B. (2005a). The Case for a new class map. *American Journal of Sociology*, 111, 141-212.

Weeden, K. A., & Grusky, D. B. (2005b.) Is inequality becoming less organized? Unpublished manuscript. Department of Sociology, Cornell University.

Weeden, K. A., & Grusky, D. B. (2005c.) Are there any big classes at all? In D. Bills (Ed.), *The shape of social inequality: stratification and ethnicity in comparative perspective. Research in Social Stratification and Mobility* (volume 22). Amsterdam: Elsevier.

Wilensky, H. L. (1966.) The professionalization of everyone? *American Journal of Sociology,* 70, 137-58.

Appendix

The analyses reported in this paper were carried out with two big-class schemes (FH and EG), two occupation schemes (indigenous and 1970-basis), and two measures of inequality (Theil and variance of log wages). Although we cannot present all possible permutations of results from our supplementary analyses, it is useful to present some of the more important ones that, for brevity's sake, were glossed over in the body of the paper.

To this end, Figures A1 and A2 present Theil decompositions of wage inequality for the male and female subsamples, decompositions that may be compared to those presented in Figures 4 and 5. The Theil index shows a more gradual increase in total wage inequality for both male and female workers than does the variance of log(wages), a difference that may be attributed, at least in part, to a correspondingly more gradual increase in within-occupation wage inequality under the Theil index. The other trendlines reveal only exceedingly slight differences and certainly do not suggest any major revisions or caveats to the substantive story outlined in the body of the paper.

In Figure A3, we have presented trends in the components of wage inequality when the big-class structure is represented by the EG scheme instead of the FH scheme and, analogously, when the micro-class structure is represented by the 1970-basis occupation scheme instead of the indigenous occupation scheme. This figure reveals, firstly, that trends in between-class and within-class inequality are virtually identical whether we measure the big-class structure with the EG or FH schemes.[16] Although class analysts are prone to disagree over measurement, at least in this case it would be difficult indeed to find reason to prefer one scheme over another.

The choice of occupation schemes, by contrast, is more consequential, although not enough so to alter in any important way the conclusions of the paper. The increase in the within-occupation component is of course more pronounced under the 1970-basis occupation scheme than under the indigenous occupation scheme. Conversely, the increase in the two remaining components of

inequality is less pronounced, so much so that the between-occupation component in 2004 is only slightly greater than the between-occupation component in 1973. These disparities are most apparent in 1983, when the BLS first adopted the 1980 SOC codes, and in 2003, when it first adopted the 2000 SOC codes. In the periods bracketed by these scheme changes, the trend lines are virtually parallel across the two occupation schemes, implying that the 1970-based scheme does not deteriorate any more rapidly in its ability to capture wages than do the indigenous occupation schemes. These cross-scheme disparities suggest that the back-coding procedure underlying the 1970-basis scheme introduces measurement error that exaggerates within-occupation heterogeneity and, conversely, suppresses the other two components of wage inequality.

Table 1.

Estimated linear effect of time on the components of wage inequality, by sex.

	Intercept	Slope coeff.	Average % change / year	% of total change
Model 1: Changing occupational structure				
A. Men				
Total	0.2369	0.0042	1.8	n/a
	(0.005)	(0.000)		
Between-class	0.0408	0.0019	4.6	45.2
	(0.002)	(0.000)		
Within-class / Between-occupation	0.0335	0.0007	2.1	16.6
	(0.002)	(0.000)		
Within-occupation	0.1626	0.0016	1.0	38.2
	(0.003)	(0.000)		
B. Women				
Total	0.1681	0.0046	2.7	n/a
	(0.005)	(0.000)		
Between-class	0.0433	0.0015	3.4	32.0
	(0.003)	(0.000)		
Within-class / Between-occupation	0.0209	0.0009	4.4	19.9
	(0.001)	(0.000)		
Within-occupation	0.1039	0.0022	2.1	48.1
	(0.002)	(0.000)		

Model 2: Constant occupational structure

A. Men

Total	0.2336	0.0026	1.1	n/a
	(0.003)	(0.000)		
Between-class	0.0403	0.0014	3.4	51.7
	(0.001)	(0.000)		
Within-class / Between-occupation	0.0346	0.0003	0.8	10.7
	(0.001)	(0.000)		
Within-occupation	0.1587	0.0010	0.6	37.5
	(0.002)	(0.000)		

B. Women

Total	0.1693	0.0024	1.4	n/a
	(0.005)	(0.001)		
Between-class	0.0449	0.0004	0.8	15.1
	(0.003)	(0.000)		
Within-class / Between-occupation	0.0220	0.0005	2.0	18.8
	(0.001)	(0.000)		
Within-occupation	0.1024	0.0016	1.5	65.7
	(0.002)	(0.000)		

Note: Numbers in parentheses are standard errors. Model 1 fits a simple linear trend on the variance in logged wages. Model 2 fits a uniform linear trend with intercept adjustments at each change in the occupational classification scheme.

Table 2.

Predicted change proportion of total variance attributable to the components of inequality, 1973 and 2004.

	Model 1		Model 2	
	1973	2004	1973	2004
A. Men				
Total variance	0.237	0.366	0.234	0.314
Percentage of total				
Between-class	17.2	27.1	17.2	26.1
Within-class / Between-occupation	14.2	15.0	14.8	13.8
Within-occupation	68.6	57.9	67.9	60.1
B. Women				
Total variance	0.168	0.311	0.169	0.243
Percentage of total				
Between-class	25.8	28.7	26.5	23.0
Within-class / Between-occupation	12.5	15.9	13.0	14.8
Within-occupation	61.8	55.5	60.5	62.2

Note: See Table 1 stub for model details.

Figure Captions

Figure 1. The transition to a big-class regime.

Figure 2. The transition to a micro-class regime.

Figure 3. Trends in men and women's wage inequality in the United States, 1973-2004.

Figure 4. Variance-based decomposition of men's wage inequality into between-class, within-class/between-occupation, and within-occupation components, 1973-2004.

Figure 5. Variance-based decomposition of women's wage inequality into between-class, within-class/between-occupation, and within-occupation components, 1973-2004.

Figure 6a. Between-occupation variance in men's logged wages in the nonmanual classes (3-year moving average).

Figure 6b. Between-occupation variance in men's logged wages in the manual classes (3-year moving average).

Figure 6c: Between-occupation variance in women's logged wages in the nonmanual classes (3-year moving average).

Figure 6d. Between-occupation variance in women's logged wages in the manual classes (3-year moving average).

Figure 7a. Within-occupation variance in men's logged wages in the nonmanual classes (3-year moving average).

Figure 7b. Within-occupation variance in men's logged wages in the manual classes (3-year moving average).

Figure 7c. Within-occupation variance in women's logged wages in the nonmanual classes (3-year moving average).

Figure 7d. Within-occupation variance in women's logged wages in the manual classes

(3-year moving average).

Figure A1. Theil-based decomposition of men's wage inequality, 1973-2004.

Figure A2. Theil-based decomposition of women's wage inequality, 1973-2004.

Figure A3. Comparison of decomposition of variance in men's (log) wages across big-class and micro-class schemes, 1973-2004.

Footnotes

[1] The corresponding rise in BC under this scenario is of course a compensating form of pro-class change.

[2] We privilege hourly earnings over weekly earnings for two reasons. First, we prefer a measure of wage inequality that is unaffected by inequality in hours worked, the latter form of inequality being conceptually distinct (and of course partially endogenous). Second, we are convinced by Lemieux's (2003) assessment that hourly earnings are subject to less measurement error than weekly earnings, given that a (slight) majority of workers in the United States prefer to report their wages in hourly terms.

[3] Although the supplements to the March CPS provide annual income for self-employed workers as well as wage and salary workers, the May/ORG files are preferable because they provide larger sample sizes, allow analysts to distinguish between the main job and all jobs, include "point-in-time" measures of wages, and rest on more consistent topcoding procedures over time (see, e.g., Card & DiNardo 2002; Lemieux 2003). Moreover, the accuracy of the March income data for self-employed workers is suspect, leading most analysts of these data to exclude these workers anyway (e.g., Card & DiNardo 2002, p. 746).

[4] We also assessed a procedure favored by Morgan & Zhang (2005; see also Mishel, Bernstein, & Schmitt 2001; Schmitt 2003) in which topcoded earnings are adjusted by year-specific multipliers calculated under the assumption that wages follow a Pareto distribution. The trends estimated under this method are very similar to those estimated under the simpler, constant-multiplier method, but with a few implausibly sharp (i.e., single-year) spikes in the observed levels of inequality. Throughout

the paper, we report results calculated from data in which topcoded values are multiplied by a constant.

[5] We have also completed analyses in which all workers were given equal weight. Not surprisingly, trends for men were unaffected, while trends for women were only trivially affected.

[6] For a given 2000 occupation, suppose that 90 percent of incumbents would have been coded into occupation X in the 1990 scheme, while 10 percent would have been coded into occupation Y. Each person with this 2000 occupation contributes two records to the expanded data set: one record receives code X and a weight of 0.9, and the other receives code Y and a weight of 0.1.

[7] We do not aggregate detailed occupations into the 126-category scheme used in much of our prior research (e.g., Weeden & Grusky 2005; Weeden 2005a). This aggregation was necessary because some of our samples, especially those based on the General Social Survey (Davis, Smith, and Marsden 2004), were quite small. We of course now have sample sizes that can support the full SOC schemes. However, just as we did not claim then that a 126-category scheme was optimal in identifying occupational boundaries, so too we wish to avoid now reifying the Census occupation schemes. We merely believe that these schemes will offer the best-available estimate of trends in inequality at the micro-class level.

[8] Morgan and Tang (2005, Supplementary Appendix pp. S2-S5), acting on a suggestion by John Goldthorpe, used earnings profiles to assign "borderline" occupations to EG codes. The resulting EG categories obviously will better capture variation in earnings than would categories devised without the benefit of earnings information. Although this coding procedure will increase the explanatory power of the EG scheme, *trends* in between-class and within-class inequality should be relatively unaffected.

[9] We also applied an alternative coding approach in which the 1970-basis EG scheme is built with pre-existing algorithms devised for the 1960 SOC codes (see, e.g., Weeden & Grusky 2005a). Where the two algorithms disagreed, we privileged the Morgan-Zhang (2005) version *except* for a few cases in which small 1970 SOC occupations were folded into large 1980 SOC occupations, yielding in the process quite implausible class assignments.

[10] We also calculated Gini coefficients for men and women. Because trends in Gini coefficients parallel trends in the Theil index, we have excluded them from Figure 3 in an effort to minimize visual clutter.

[11] We have also calculated the analogous percentages based on observed values of BC, BO/WC, and WO in 1973 and 2004. The results are (unsurprisingly) similar to those secured under Model 1.

[12] The suppressed BC slope coefficients for women may be understood as an artifact of the unusually steep downward trend in the BC series in the 1970s. If one fits a linear trend to the post-1980 data, the gender difference in the slope withers away.

[13] This conclusion does not hold for the "Model 2" specification for men.

[14] The "Model 2" specification for women is again inconsistent with this conclusion.

[15] The declining wages in food service occupations is no doubt partly attributable to the declining minimum wage (in real dollars).

[16] This non-finding is perhaps surprising, given that our prior research has shown that the FH scheme can better capture inequalities in most outcomes, including income (Weeden & Grusky, 2005a,b). The earlier finding was, however, based on an implementation of the EG scheme that did not incorporate earnings profiles as a criterion for class assignment (see footnote 8).

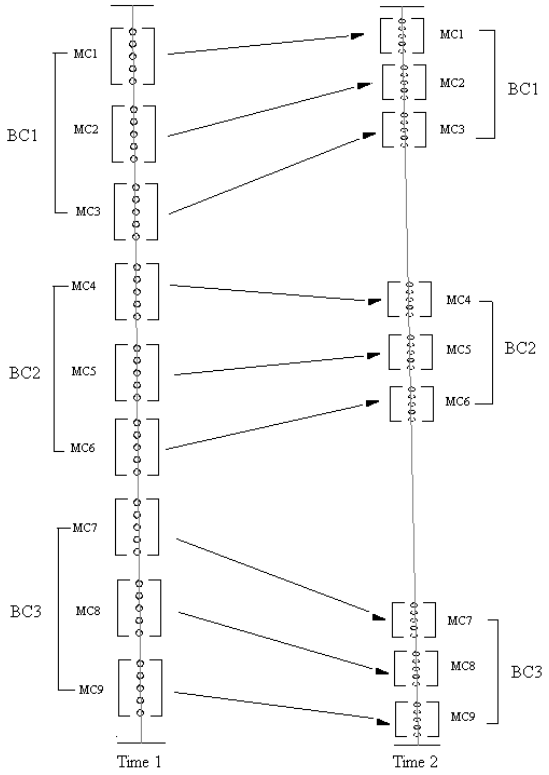

Figure 1. The Transition to a big-class regime

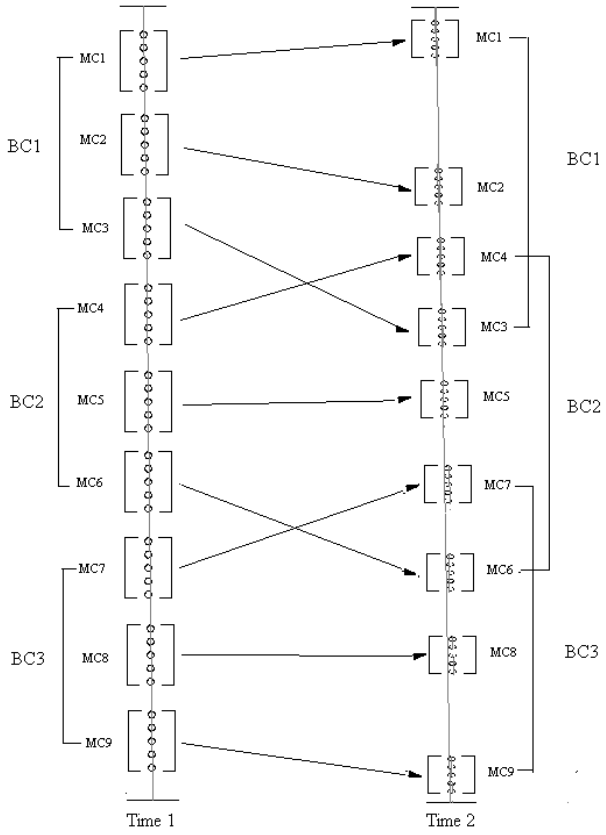

Figure 2. The Transition to a micro-class regime

Figure 3: Trends in men and women's wage inequality in the United States, 1973-2004

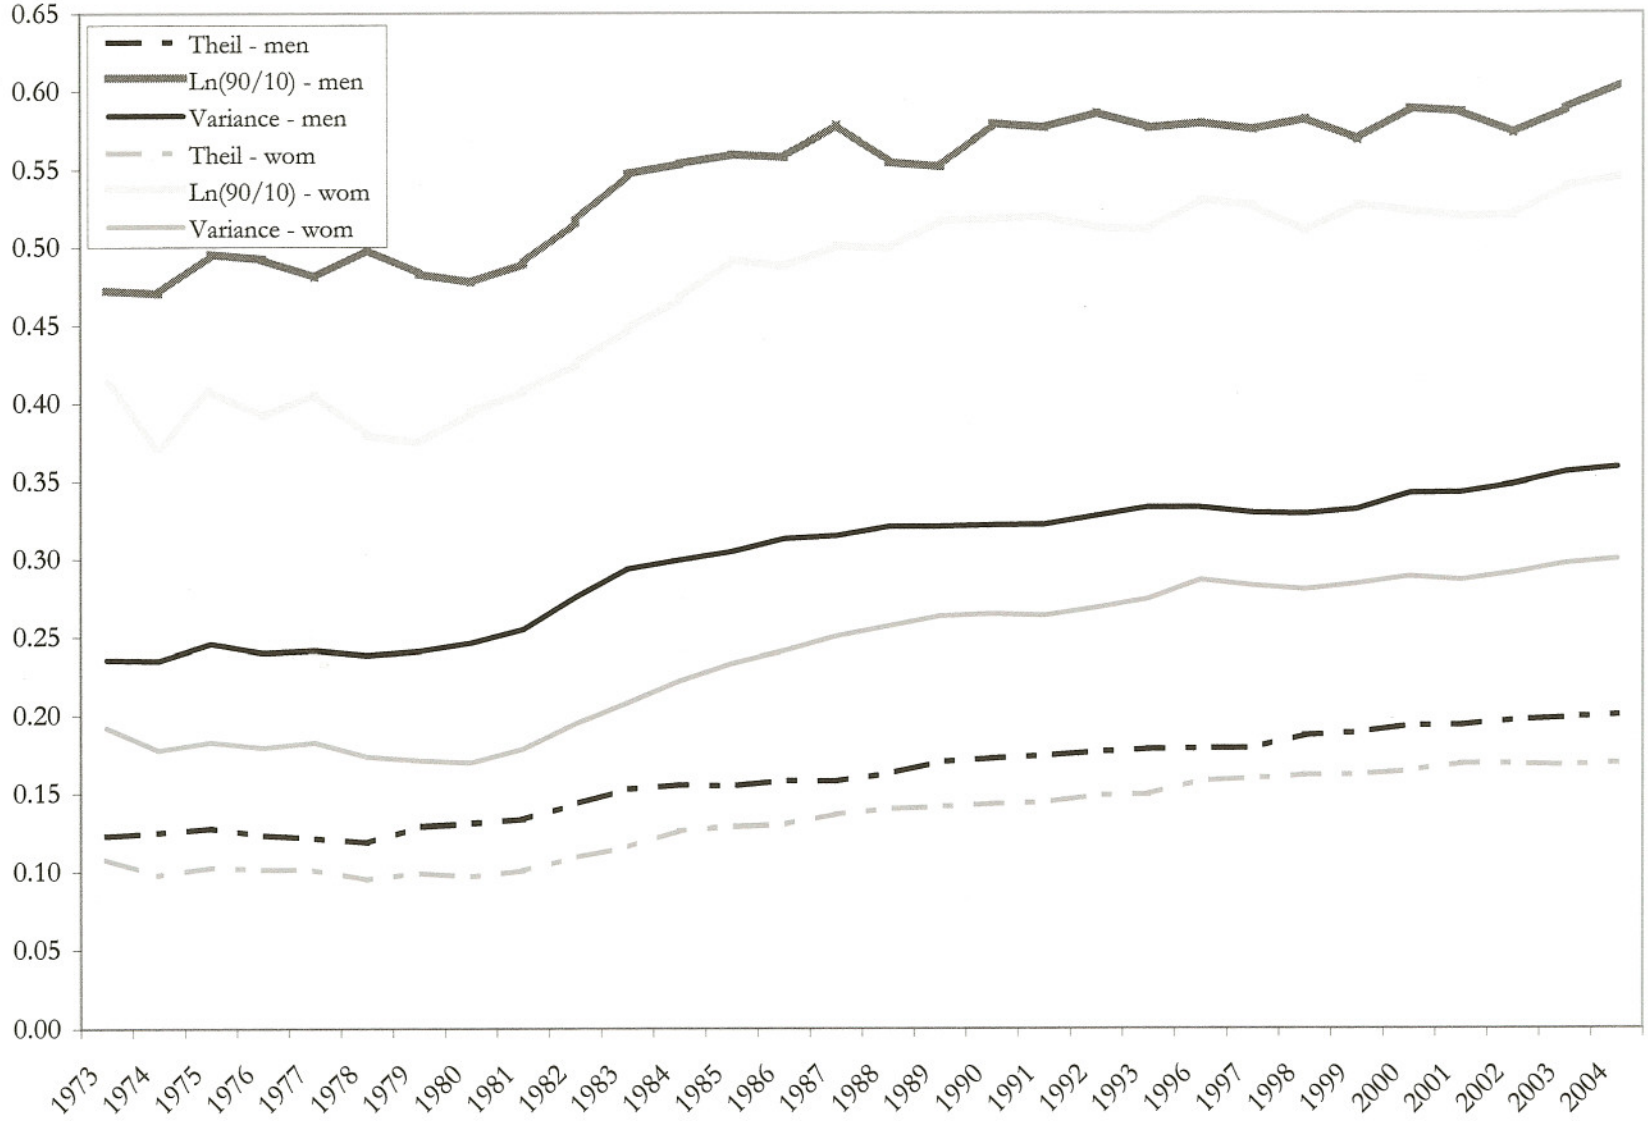

Figure 4: Variance-based decomposition of men's wage inequality into between-class, within-class/between-occupation, and within-occupation components, 1973-2004

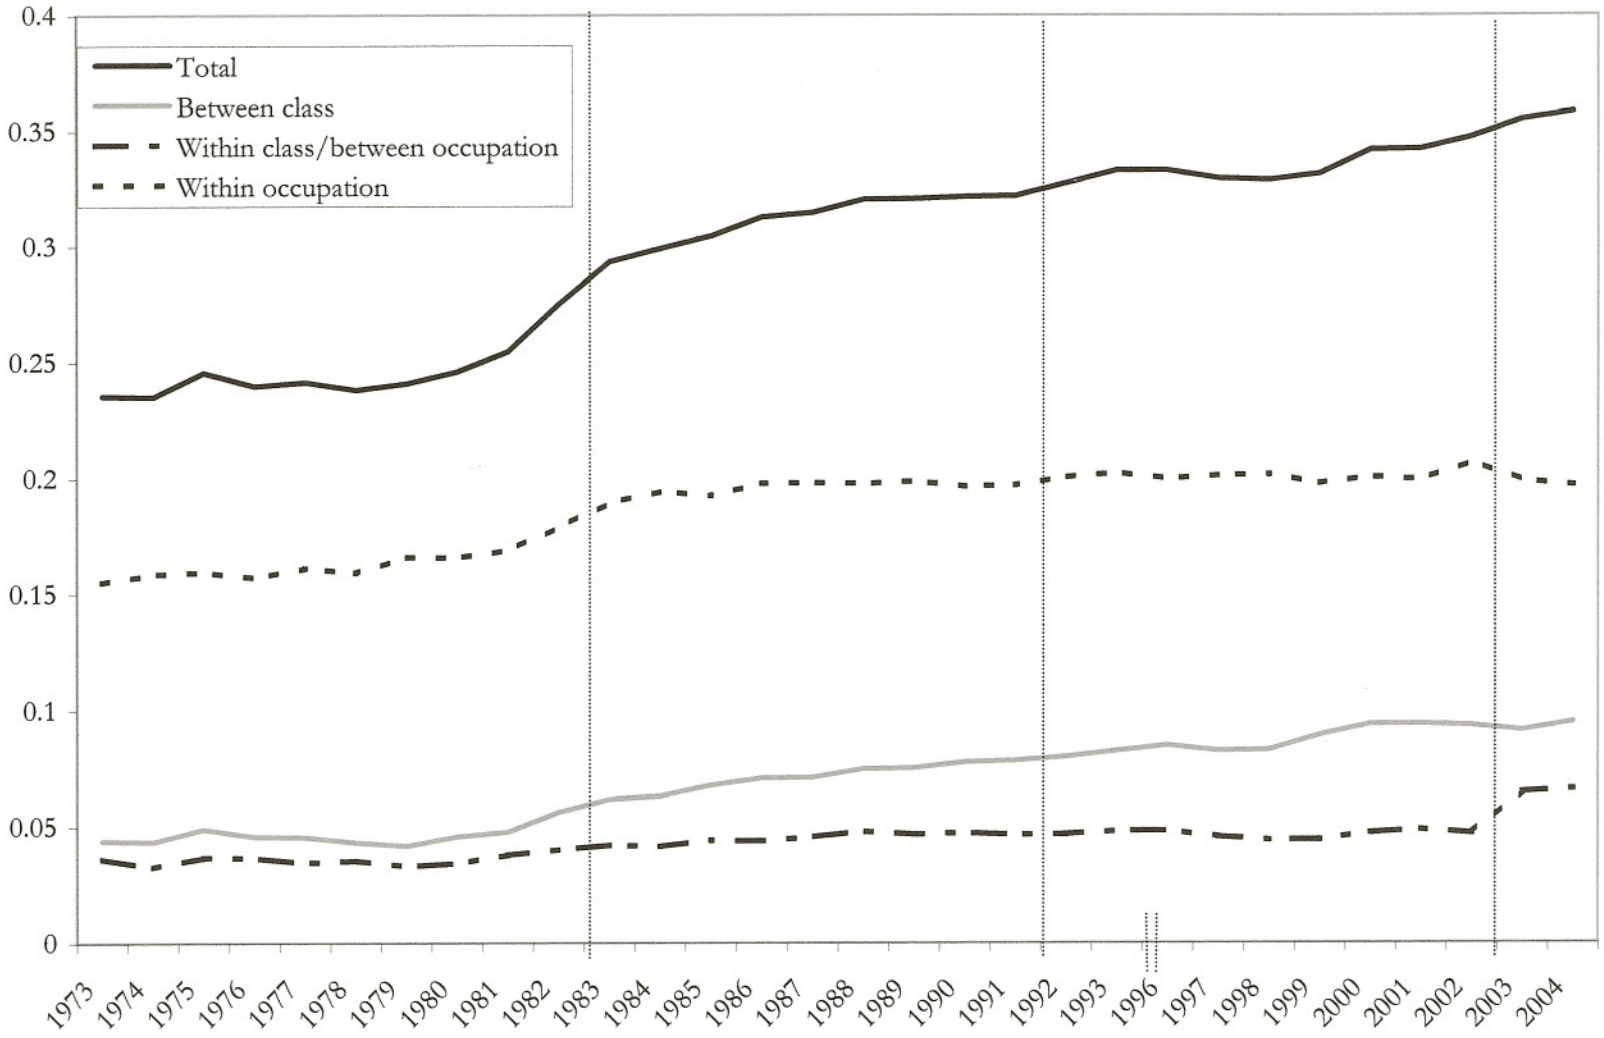

Note: Vertical lines indicate new SOC schemes, hash marks indicate break in time series

Figure 5: Variance-based decomposition of women's wage inequality into between-class, within-class/between-occupation, and within-occupation components, 1973-2004

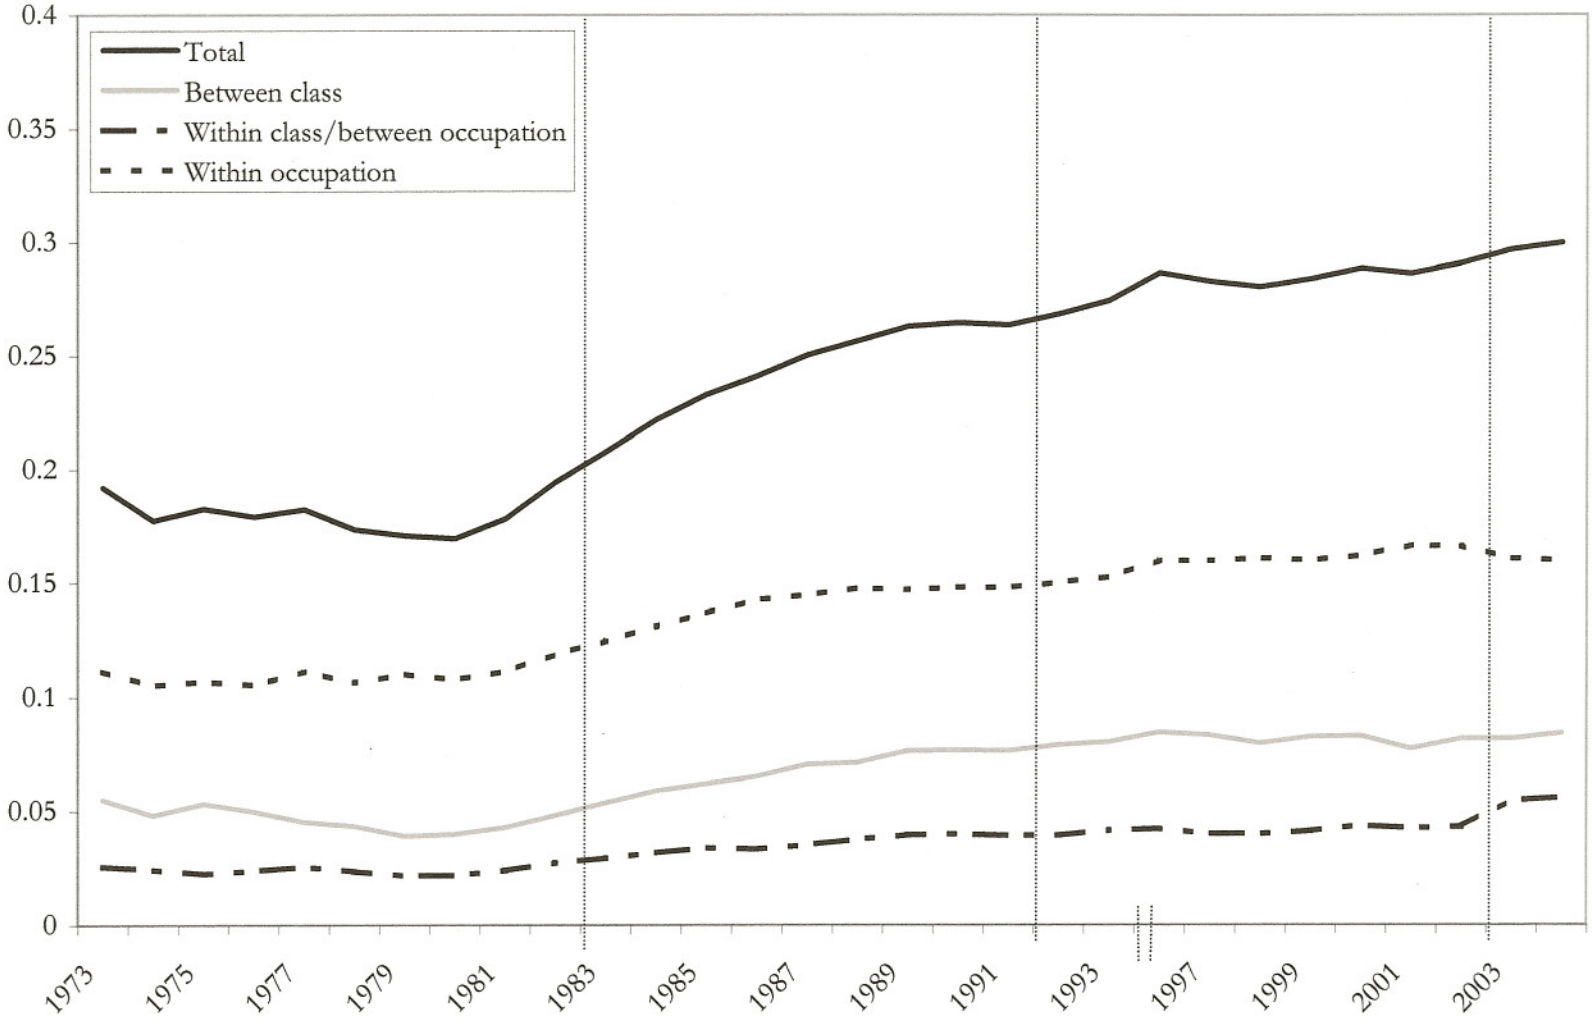

Note: Vertical lines indicate new SOC schemes, hash marks indicate break in time series

Figure 6a: Between-occupation variance in men's logged wages in the nonmanual classes
(3-year moving average)

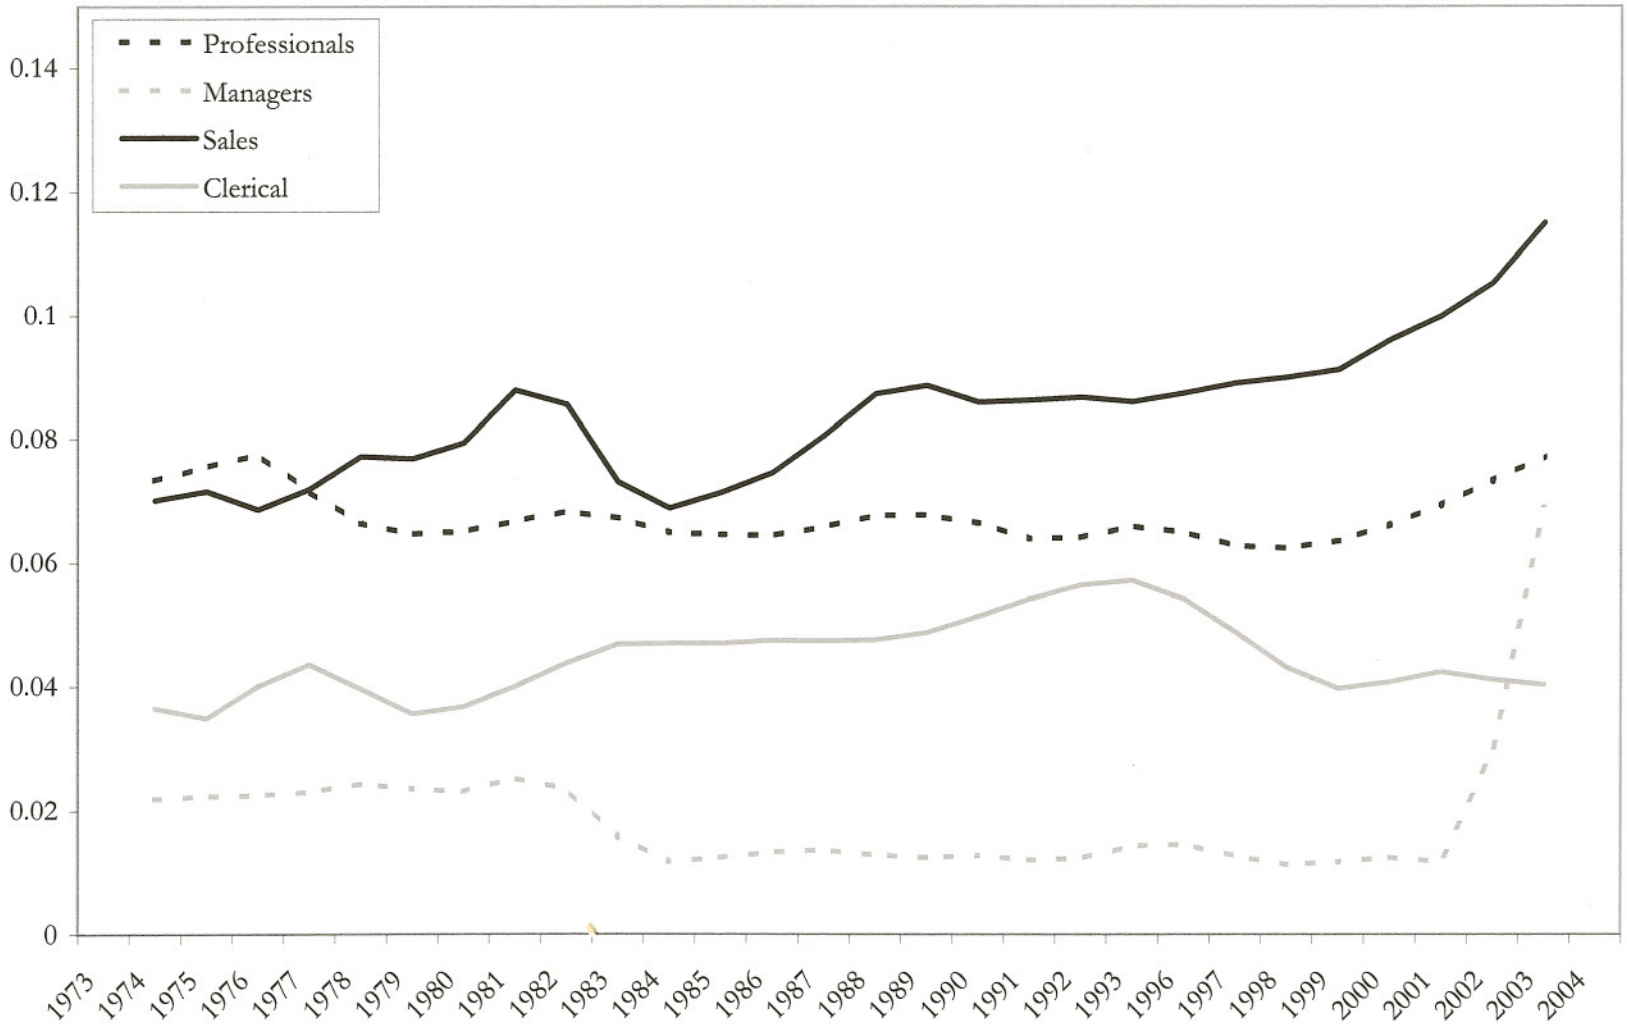

Figure 6b: Between-occupation variance in men's logged wages in the manual classes
(3-year moving average)

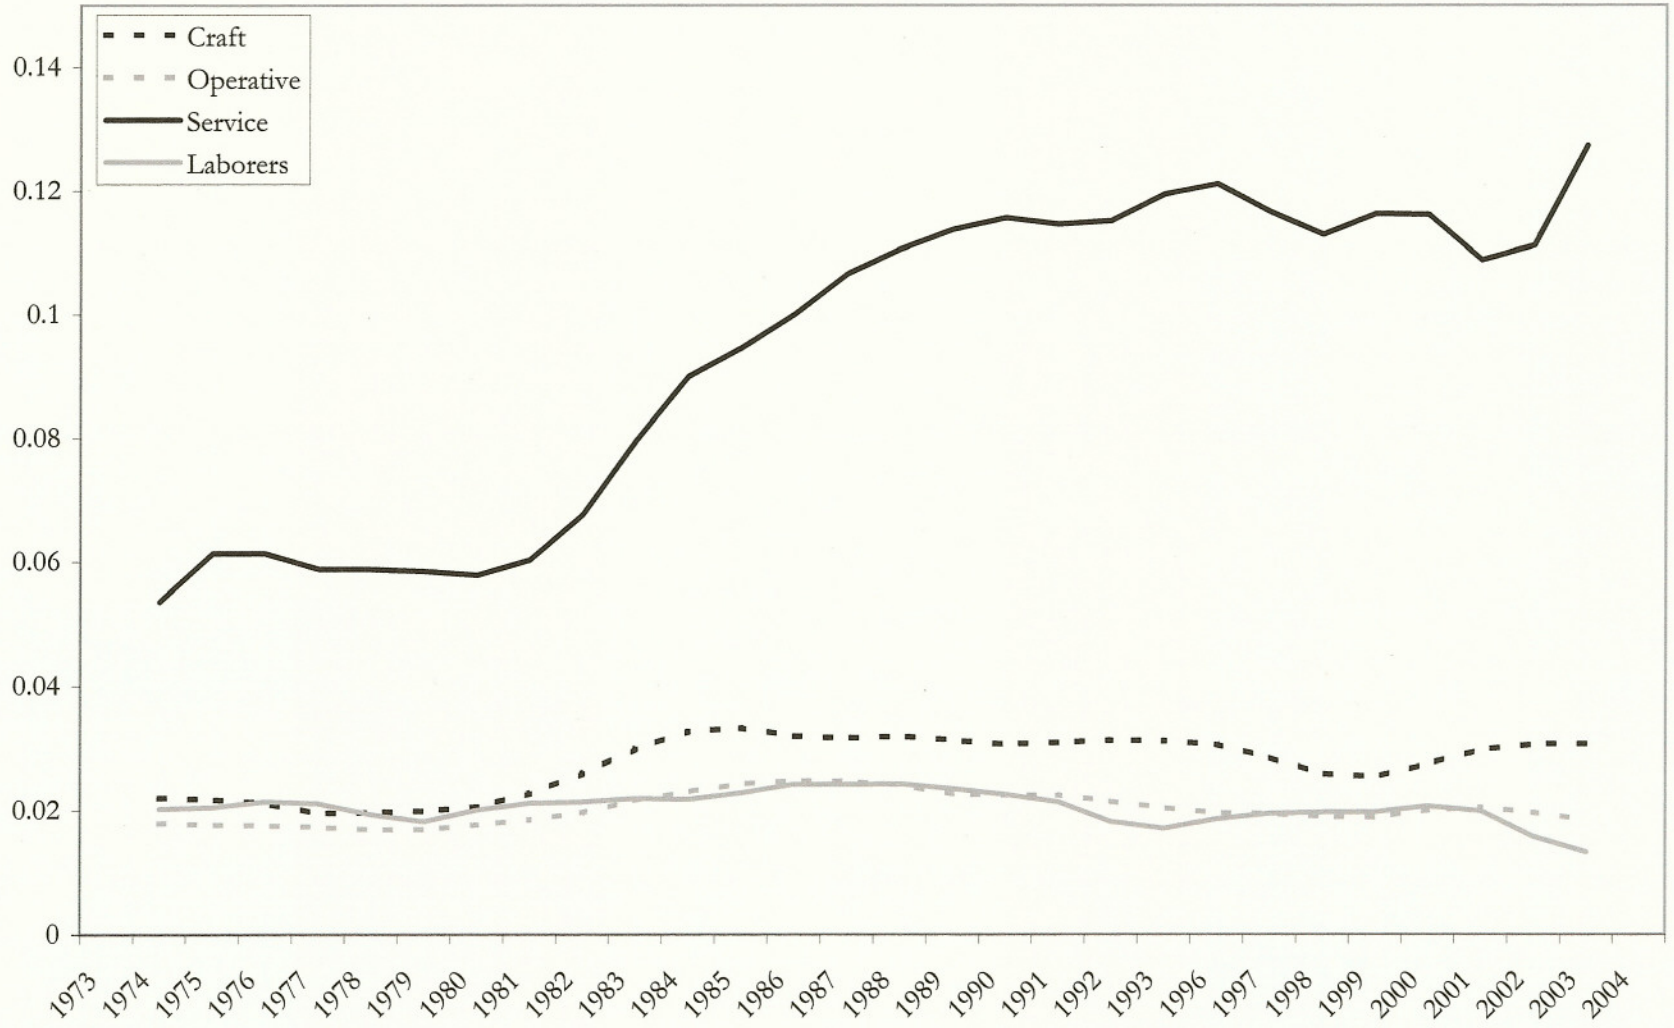

Figure 6c: Between-occupation variance in women's logged wages in the nonmanual classes
(3-year moving average)

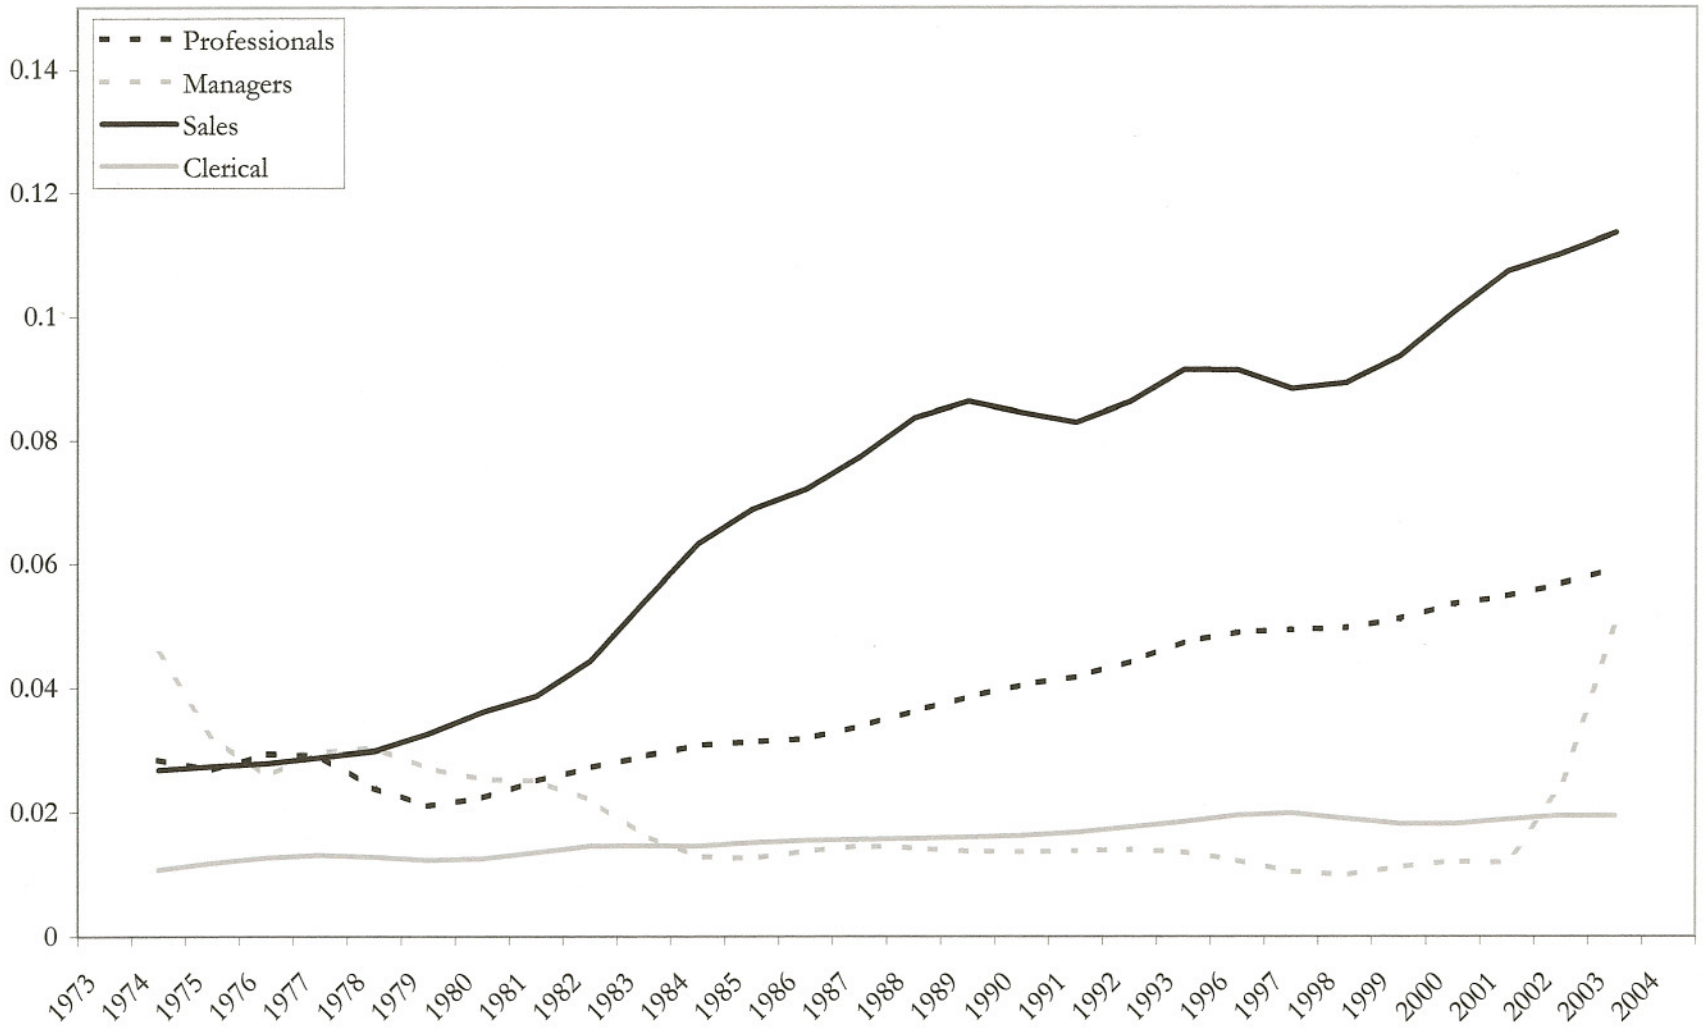

Figure 6d: Between-occupation variance in women's logged wages in the manual classes
(3-year moving average)

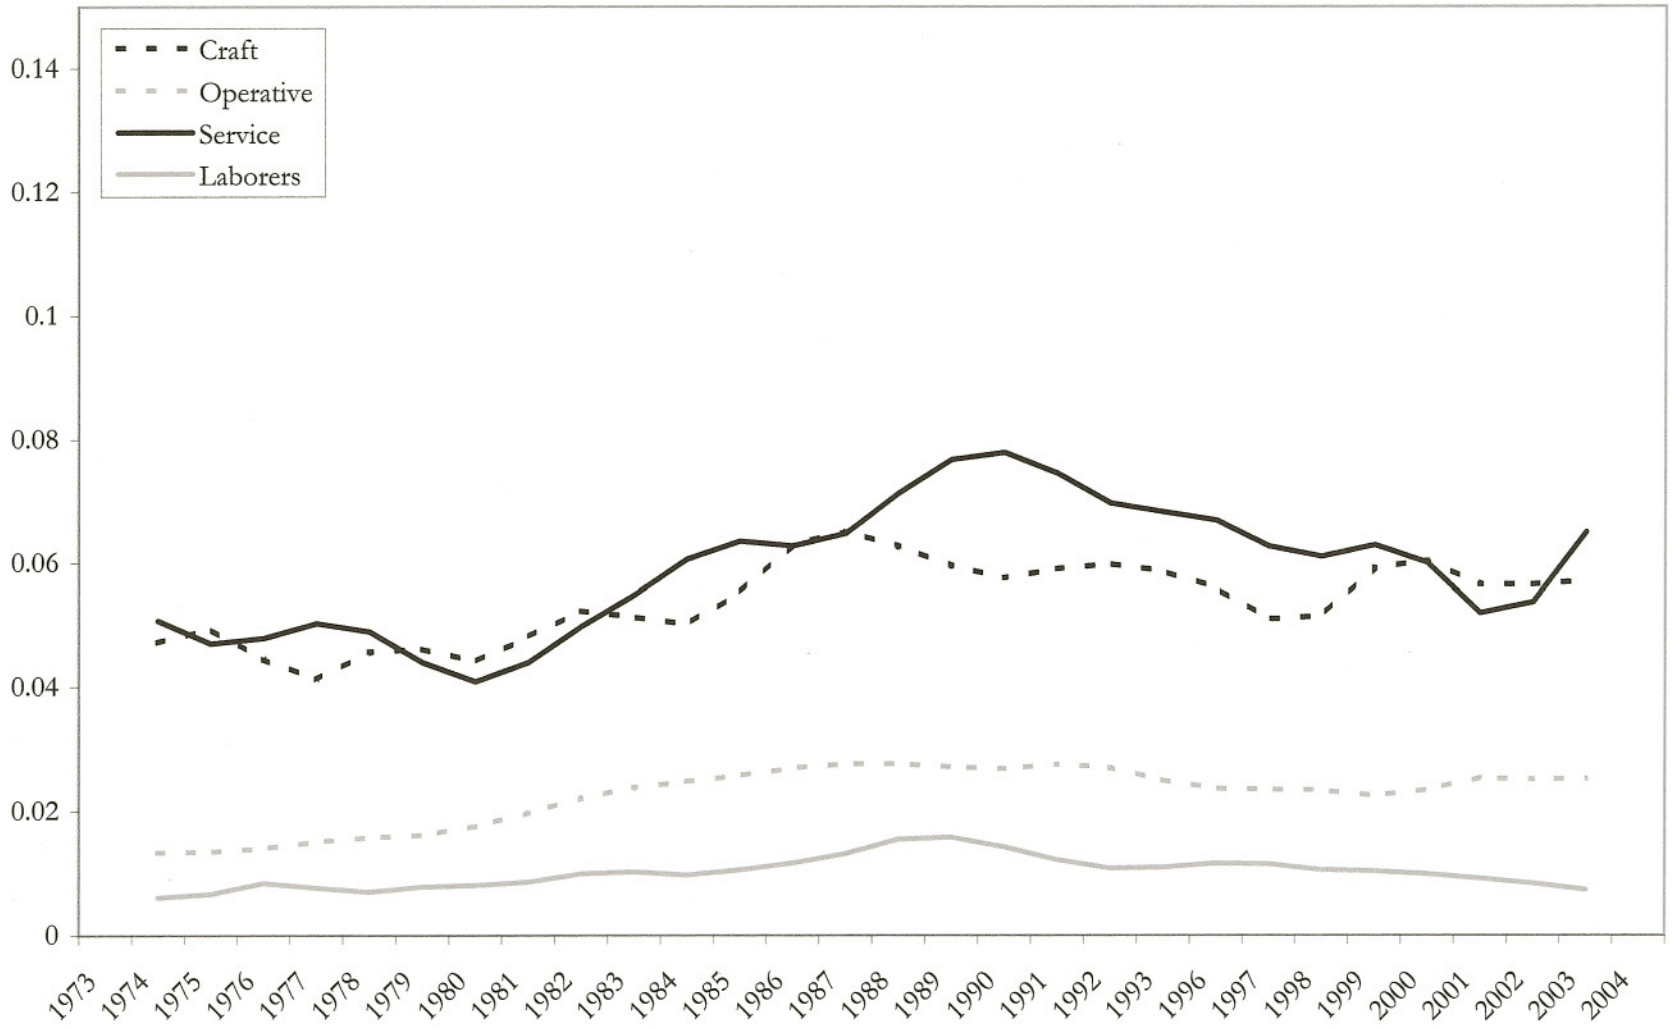

Figure 7a: Within-occupation variance in men's logged wages in the nonmanual classes
(3-year moving average)

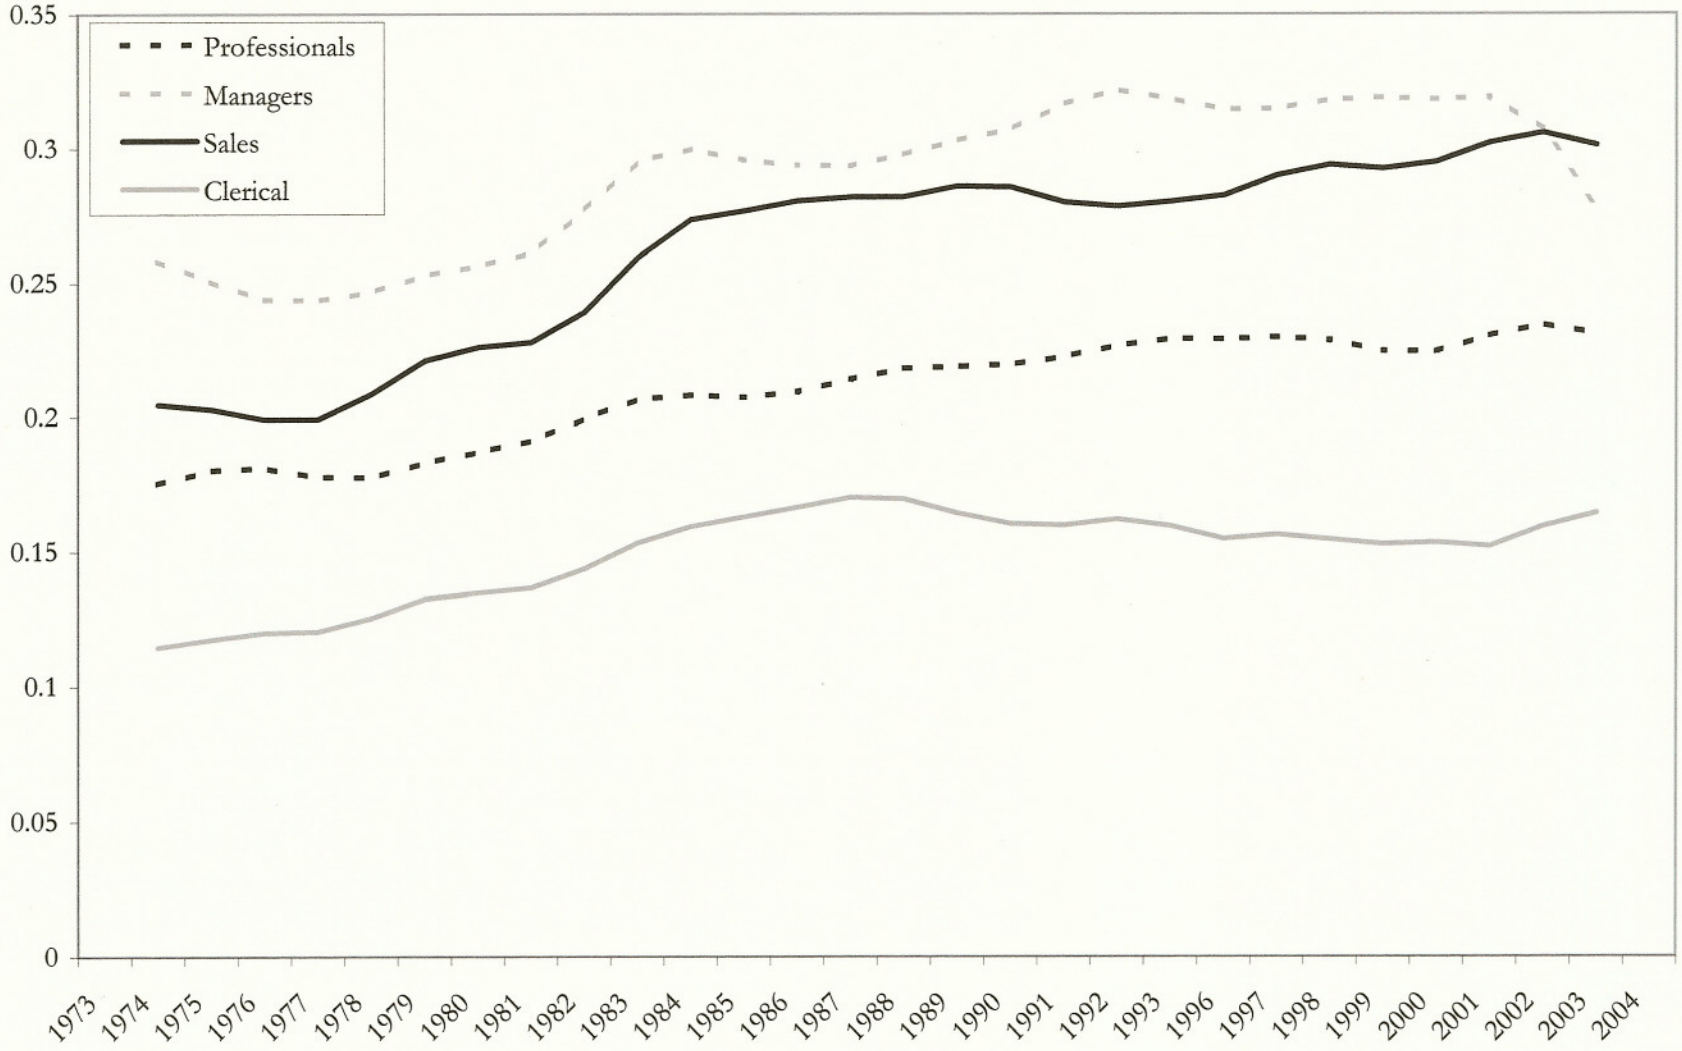

Figure 7b: Within-occupation variance in men's logged wages in the manual classes
(3-year moving average)

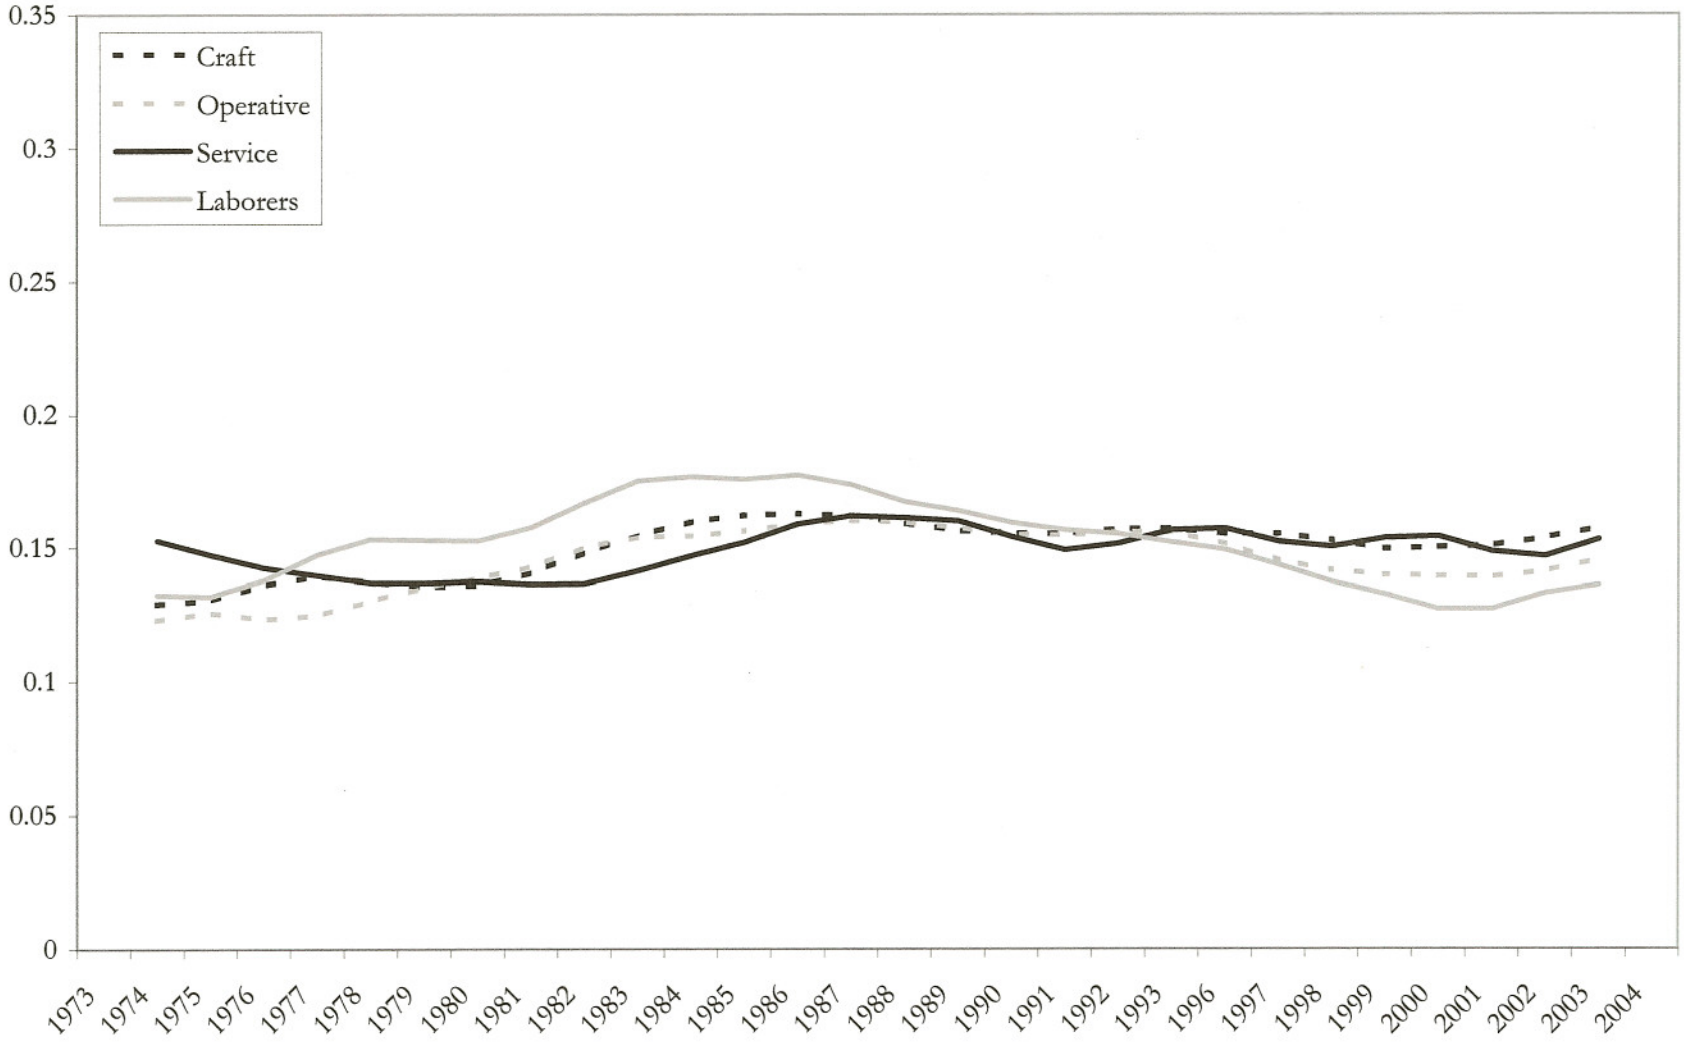

Figure 7c: Within-occupation variance in women's logged wages in the nonmanual classes
(3-year moving average)

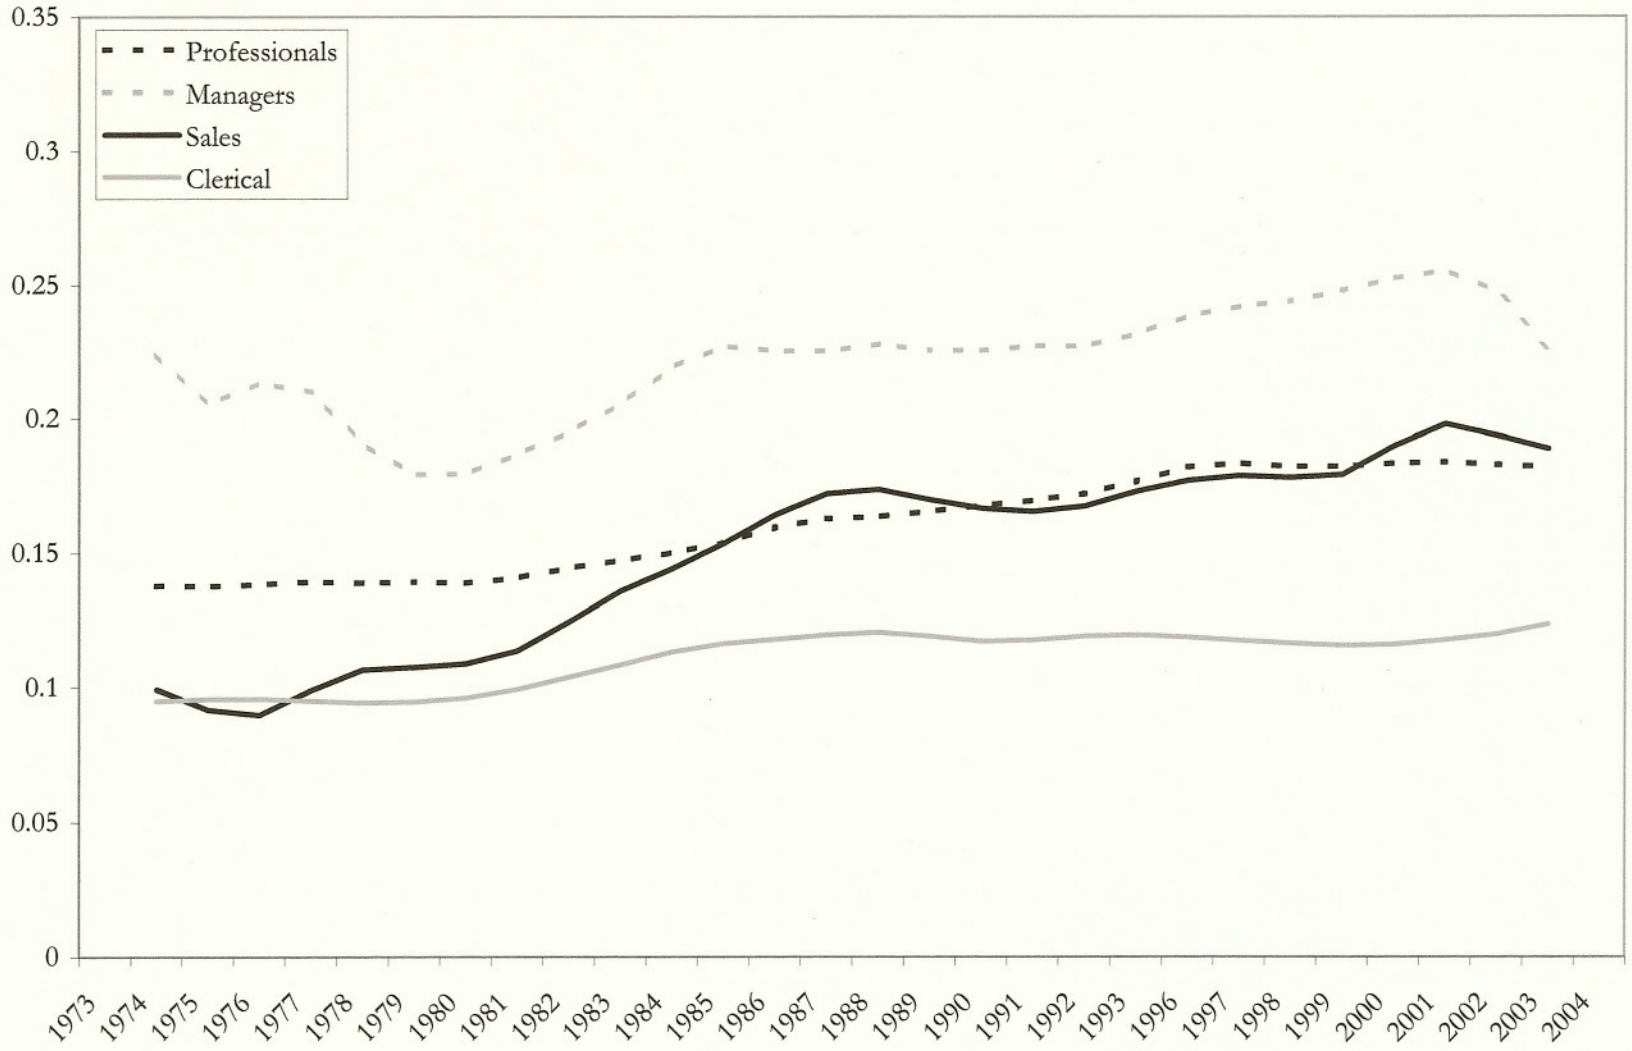

Figure 7d: Within-occupation variance in women's logged wages in the manual classes
(3-year moving average)

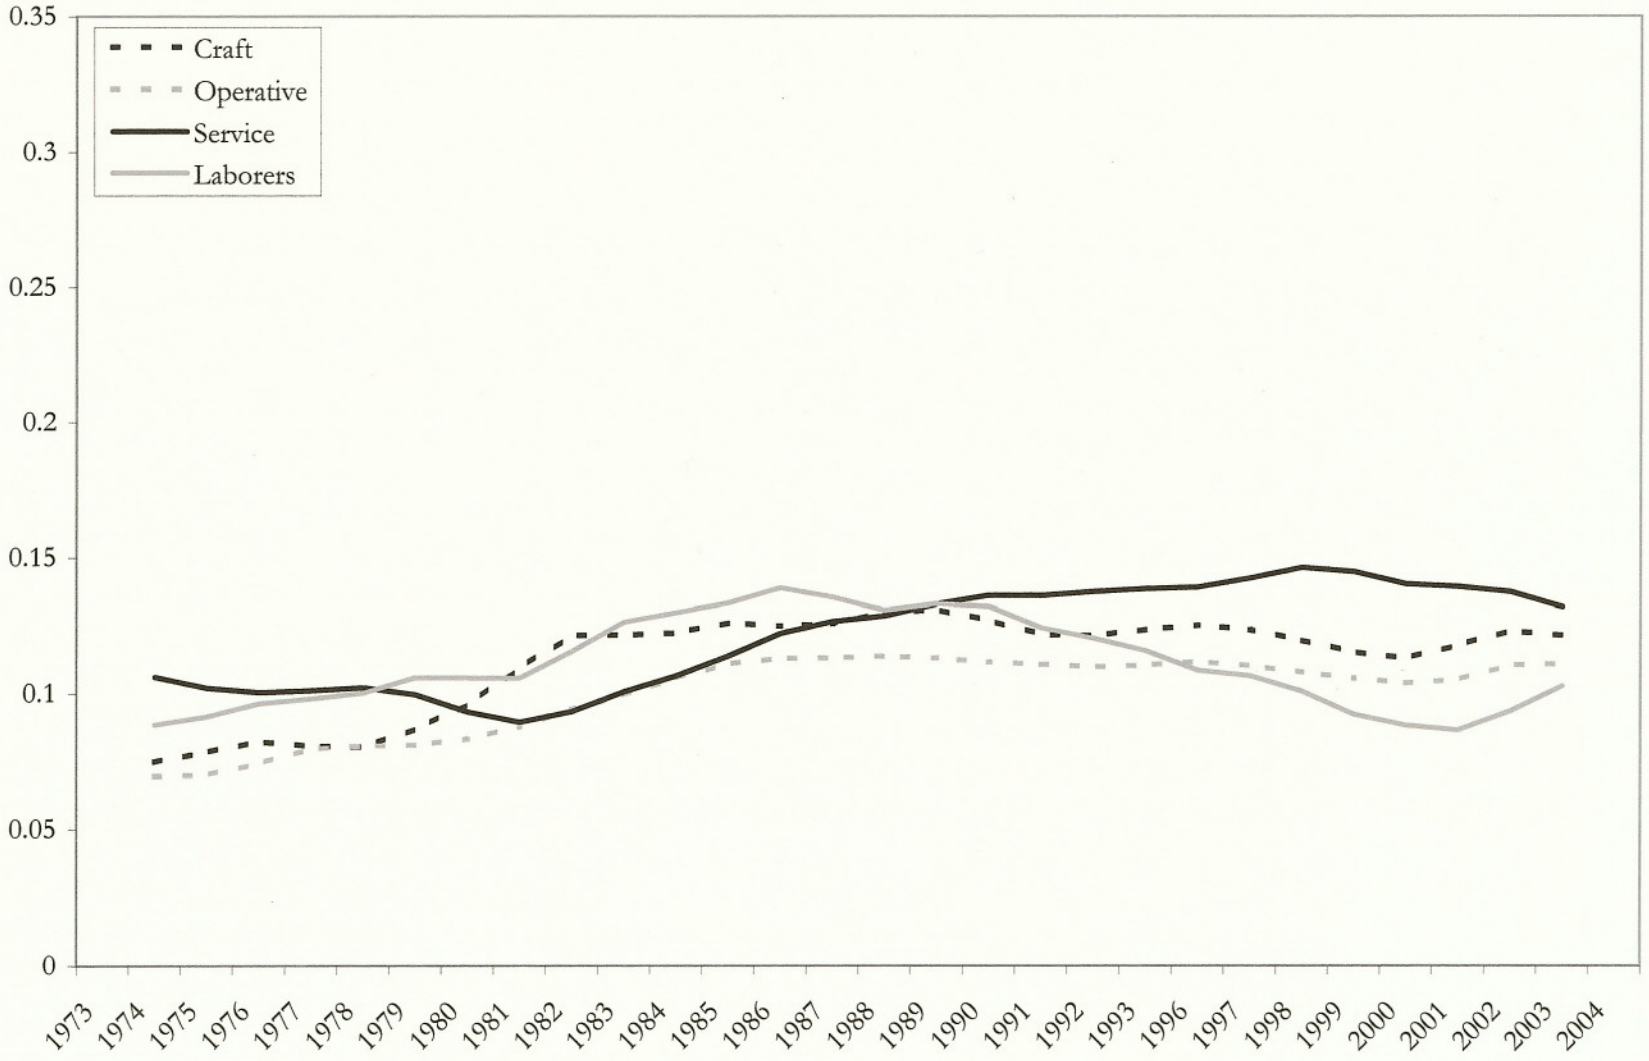

Figure A1: Theil-based decomposition of men's wage inequality, 1973-2004

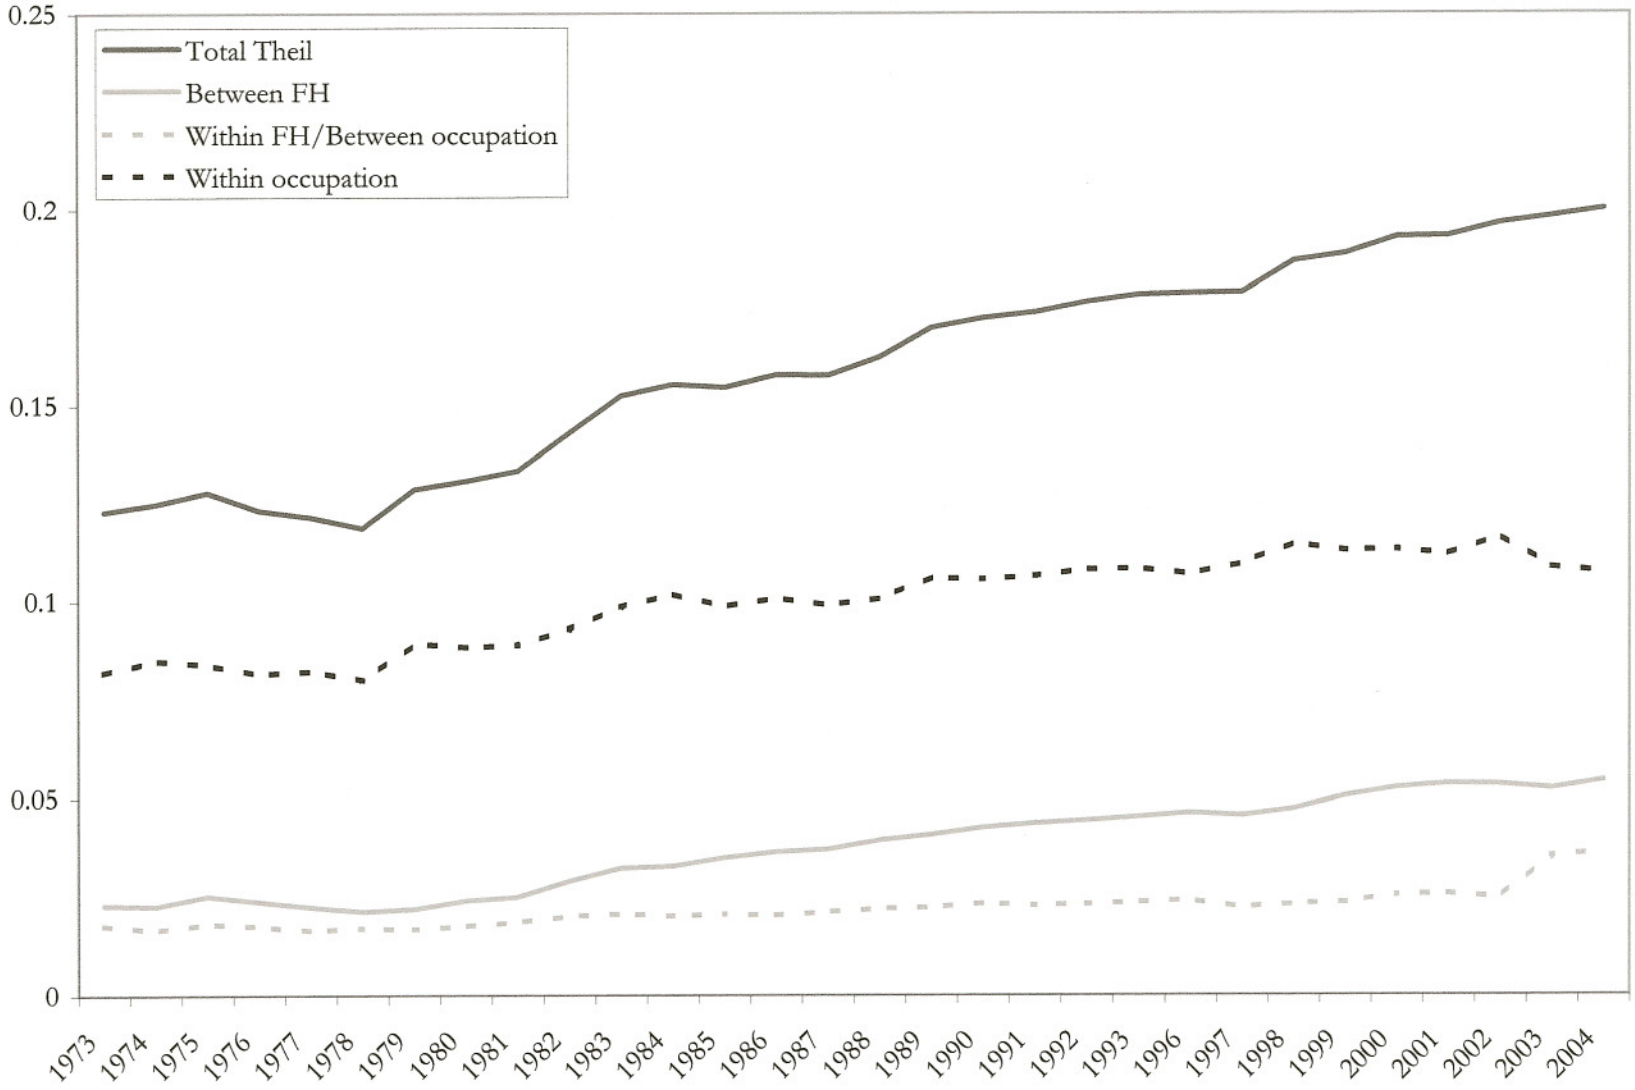

Figure A2: Theil-based decomposition of women's wage inequality, 1973-2004

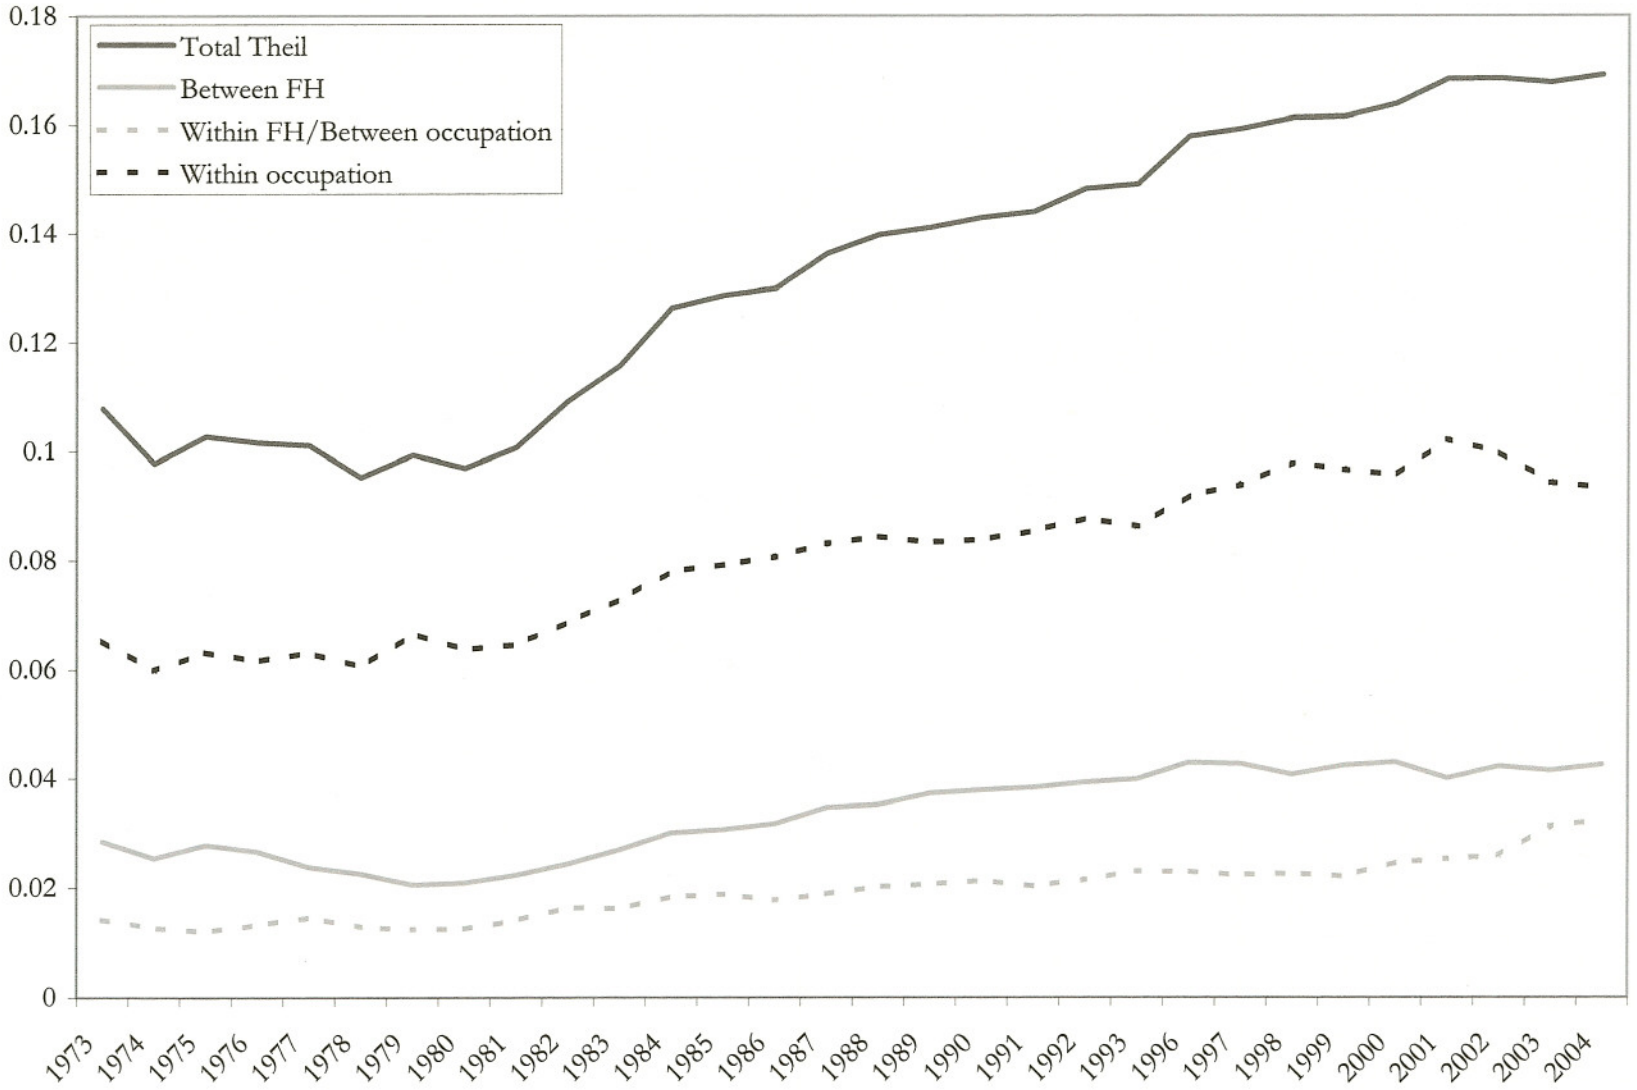

Figure A3: Comparison of decomposition of variance in men's (log) wages across big-class and micro-class schemes, 1973-2004

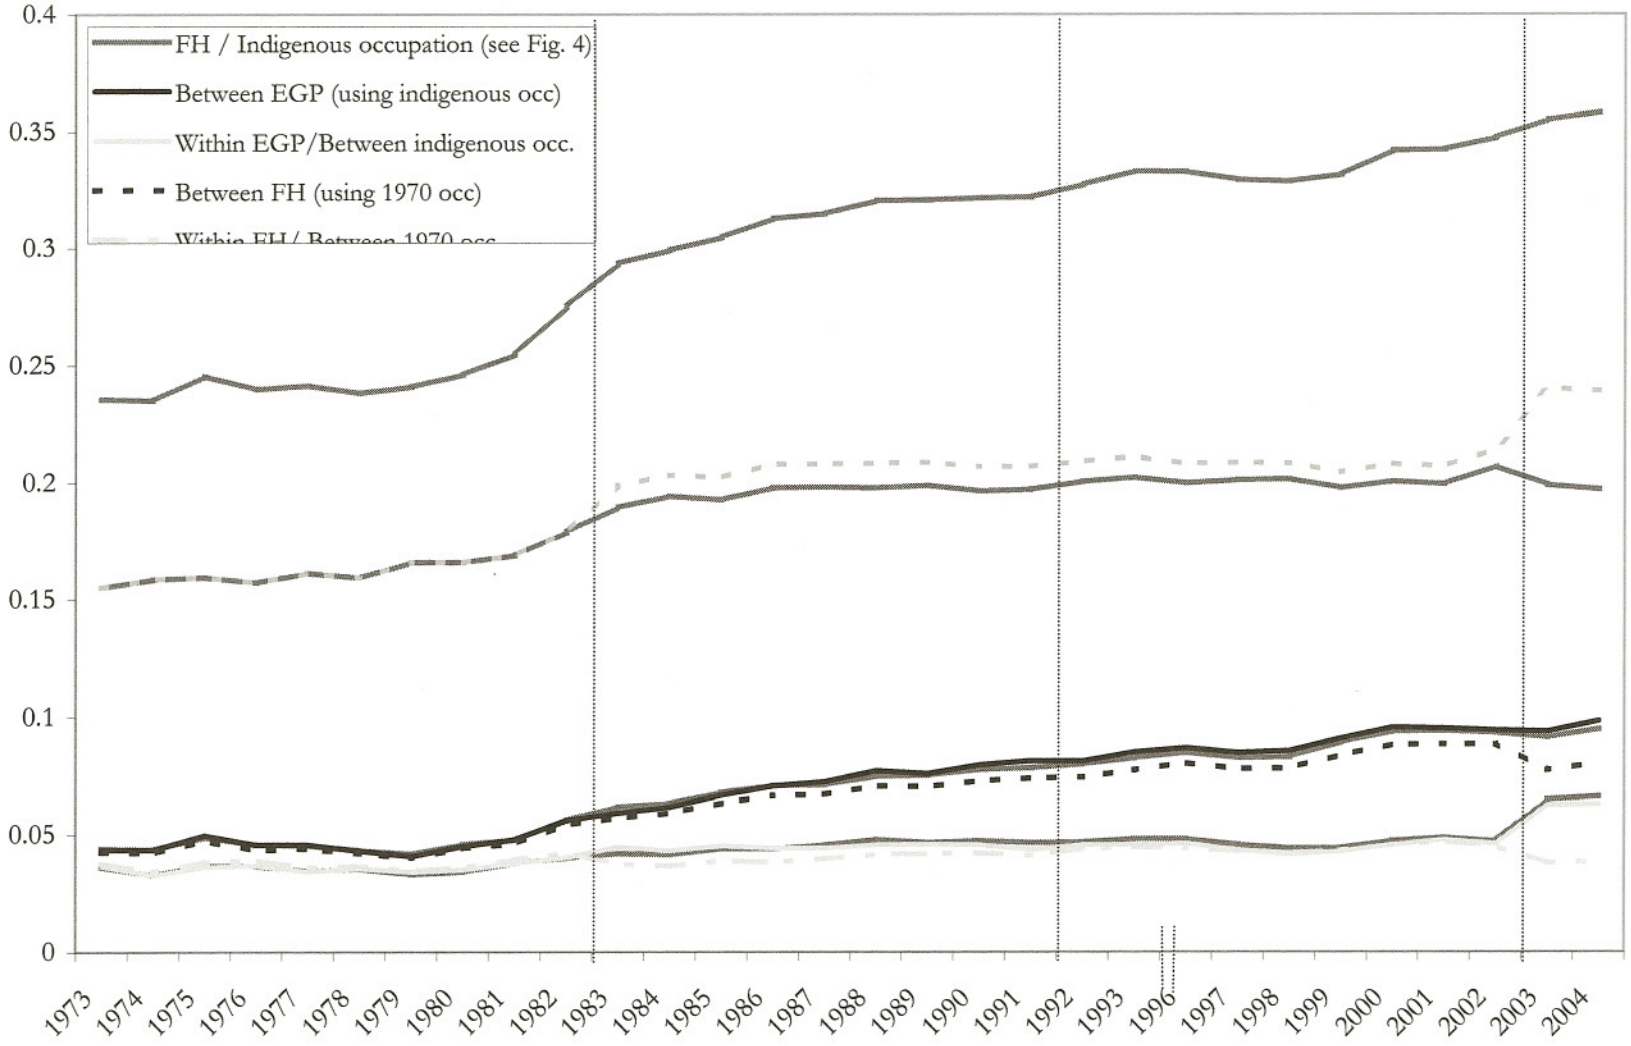

Note: Vertical lines indicate new SOC schemes, hash marks indicate break in time series