

Firms and Skills: The Evolution of Worker Sorting*

Christina Håkanson[†] Erik Lindqvist[‡] Jonas Vlachos[§]

September 28, 2012

PRELIMINARY AND INCOMPLETE

Abstract

Are workers increasingly being sorted into firms according to their skills? We investigate this issue using data on cognitive and non-cognitive skills from the military enlistment for 28 cohorts of Swedish men matched to their employers. The enlistment skill measures are comparable over time and measured before men enter into the labor market. We document a significant increase in sorting by both cognitive and non-cognitive skill from 1986 to 2008: Skill differences within firms have fallen across all major industries while differences in skill between firms have increased. The evidence point to technological change rather than outsourcing or trade as the main factor behind these changes: Between-firm skill differences increase mainly due to a reallocation of engineers following the expansion of the IT industry. Skill differences within firms fall because of stronger assortative matching of workers and not because firms perform a more narrow set of tasks.

Keywords: Skill sorting; skilled-biased technological change; outsourcing; globalization; cognitive skills; non-cognitive skills; personality; employer-employee matched data.

JEL codes: J24, J62, L21, O33

*We are thankful to David Cesarini, Jim Heckman, Fredrik Heyman, John Van Reenen, Valérie Smeets, Yoichi Sugita, Frederic Warzynski, seminar participants at Bena (Berlin), Gothenburg University, IFAU, the 2012 AEA meetings, EALE 2012 and the National Meeting of Swedish Economists 2012 for valuable comments. Financial support from the Jan Wallander and Tom Hedelius Foundation, IFAU, and Riksbankens Jubileumsfond is gratefully acknowledged.

[†]Institute for International Economic Studies, Stockholm University, Sweden. E-mail: christina.hakanson@iies.su.se

[‡]Stockholm School of Economics and Research Institute of Industrial Economics (IFN). E-mail: erik.lindqvist@hhs.se

[§]Department of Economics, Stockholm University, and Research Institute of Industrial Economics (IFN), and CEPR. E-mail: jonas.vlachos@ne.su.se.

1 Introduction

How has the sorting of workers by skill into firms evolved over recent decades? Previous empirical studies on this question have been constrained by a lack of measures of skill that are comparable over time. In this paper, we use rich data on skills for a large, representative sample of Swedish men matched to their employers in order to quantify changes in sorting over time. Our main result is that differences in skills between firms have increased since the mid 1980's while workers have become more similar within firms.

There are a number of reasons to believe that technological change and globalization increase the sorting of workers by skill to firms. For example, the theoretical literature has stressed that firms investing in skill-biased technologies face a higher return to hiring skilled workers (Acemoglu, 1999; Caselli, 1999). A different possibility is that more complex production processes strengthen the complementarity between workers skills, implying that unskilled workers constitute "weak links" in firms with skilled workers (Kremer, 1993). Globalization increases the scope for skill-sorting by narrowing the set of tasks that needs to be performed domestically (Feenstra and Hanson, 1996; Grossman and Rossi-Hansberg, 2008) and by allowing skilled workers in rich countries to match with workers in developing countries rather than unskilled workers in their own country (Kremer and Maskin 2006). Relatedly, Grossman and Maggi (2000) argue that lower trade costs induce countries with a comparably low dispersion of skill, such as Sweden, to specialize in industries where it is optimal to match workers with similar skills. The upshot of all of these proposed mechanisms is that firms become more heterogeneous in terms of the skill level of their workforces. In other words, the economy might to an increasing extent be divided into Google-type firms that employ the most able workers and firms like McDonald's that employ the least able.

An increase in the segregation by skill is likely to have substantial economic and social consequences. First, wage inequality is increasing in skill segregation if worker skills are complements (e.g. Sattinger, 1975), or if fair wage considerations compress wage differences between low- and high-skilled workers within firms (Akerlof and Yellen, 1990; Bewley, 1999). Second, the extent of social interaction between different strata in society is reduced as workplaces become more internally homogeneous, with potentially far-reaching consequences for the formation of social networks and for social cohesion in general.¹

Recent empirical evidence on labor markets in advanced economies is consistent with an increase in sorting. The increase in wage inequality witnessed in many advanced economies over recent decades appears to be associated with increasing wage differences between, rather than within, firms.² Labor markets in advanced economies are also becoming increasingly polarized as routine jobs disappear while both high- and low-skilled non-routine jobs become more prevalent (Acemoglu and Autor, 2011; Adermon and Gustavsson, 2011). Yet, relatively little is still known whether these changes in the structure of wages and jobs have been accompanied by an increase

¹See Jackson (2010) for an overview of social networks and their impact on economic behaviour.

²This literature includes studies of Sweden plants 1986-2000 (Nordström-Skans, Edin and Holmlund 2010); Czech firms 1998-2006 (Eriksson, Pytlikova and Warzynski, 2009); the US manufacturing sector 1975-1992 (Dunne et al, 2004) and UK firms 1984-2001 (Faggio et al, 2010).

in segregation of workers by skill across firms.

The difficulty in assessing changes in sorting stems from a lack of skill measures that are comparable over time. Previous research on skill sorting has either focused on educational attainment (Kremer and Maskin, 1996), occupation (Kremer and Maskin 1996; Dunne et al 1997, 2004), or skill measures derived from wage data (Iranzo, Schivardi and Tosetti, 2008). Each approach faces potential problems. For example, skilled-biased technological change may increase the dispersion of wages, even though the underlying distribution of skills remains unchanged. Relatedly, changes in the occupational structure reflect changes in technology rather than changes in the composition of skills. The level of educational attainment might also not be comparable over time: Higher education has expanded in most countries and students' choices between different fields of education change in response to the economic environment.³ Further, educational attainment, by construction, does not capture heterogeneity in skill within educational groups.⁴

In this paper, we study the evolution of skill sorting in the Swedish private sector between 1986 and 2008 using data on workers' cognitive and non-cognitive abilities from the military enlistment. The enlistment skill measures are strong predictors of future labor markets outcomes (Lindqvist and Vestman, 2011), comparable over time, and available for 28 cohorts of Swedish men. Since the enlistment evaluations were administered to Swedish men at the age of 18, the skill measures are unaffected by the expansion of higher education and changes in labor market conditions.

Matching the enlistment skill measures for workers with information about their employer in a given year, we document a substantial increase in skill segregation from 1986 to 2008. During this period, workers became more similar within firms (falling within-firm variance of skills) and more dissimilar between firms (increasing between-firm variance) with respect to both cognitive and non-cognitive skills. Moreover, we find that the between-firm covariance of cognitive and non-cognitive skill is positive and increasing over our study period. The trend toward smaller differences in cognitive skills within firms is strongest in manufacturing, where within-firm skill differences were large in 1986. Yet the shift toward smaller skill differences within firms is present in all major industries, including service industries. The trend toward more segregation of skill is robust to assuming alternative distributions of skills and non-parametric ways of measuring sorting. Using data on male relatives to impute cognitive and non-cognitive skills for women, we find the same trend toward segregation among female workers.

Three sets of results point to a key role for technological factors in explaining the increase in between-firm differences in cognitive skill over time. First, the growth of the IT industry can alone account for almost the entire increase in cognitive skill differences between industries. Second, removing workers with a technical education from the sample takes out most of the increase in sorting between firms. Third, we find that increasing differences in the skill-intensity

³Skill levels can change quite rapidly within fields of education: Grönqvist and Vlachos (2008) document that the average cognitive ability among entering teachers declined by more than half a standard deviation between 1992 and 2007.

⁴That income inequality within educational groups has increased suggests that within-group skill heterogeneities are becoming increasingly important (Machin, 1996; Katz and Autor, 1999). Altonji et al (2012) provide an overview of the returns to secondary and post-secondary education across different majors.

of technology across firms and stronger matching of skilled workers to skill-intensive firms explain the increase in between-firm skill differences. Taken together, the pattern in the data is broadly consistent with the predictions from the models by Caselli (1999) and Acemoglu (1999): After the introduction of a new technology (IT), workers with high and low cognitive skills select into different sectors. Overall, the results are similar for non-cognitive skills, but technological differences play less of a role both in the cross section and for changes over time.

While technological differences across firms increase, the trend toward decreased within-firm skill variance is driven by stronger assortative matching of skills and not by changes in the production technology within firms. In other words, increased skill homogeneity within firms is due to a lower dispersion of skill among workers performing similar tasks, rather than firms performing a more narrow set of tasks. As we expect firms that outsource non-core activities to become more homogeneous with respect to the skill requirements of the tasks that remain within the firm, this finding indicates that outsourcing is not a main driver of sorting. Nor do we find that industries with initially homogeneous firms experienced stronger growth, a finding which contradicts the conjecture of Grossman and Maggi (2000) that trade induces specialization in industries that exhibit positive or negative skill complementarities. Instead, we argue that the evidence is consistent with a fall in search costs or skill complementarities becoming more important over time.

We discuss the theoretical and empirical literature on skill sorting in more detail in the next section. We discuss our skill measures and the construction of the data set is discussed in Section 3, our approach for measuring sorting in Section 4 and the main results in Section 5. We present the corresponding sorting patterns for education and skills derived from wage data in Section 6. Section 7 analyzes the mechanisms behind the observed changes in sorting documented in Section 5. Section 8 concludes the paper. We present additional material in a set of appendices denoted A (data description), B (additional results) and C (issues regarding how to quantify sorting).

2 Literature

The optimal allocation of skill across firms depends on the nature of the production function. Changes in sorting by skill is therefore either due to changes in the production function itself, or to changes in the constraints in the matching of workers to firms. With respect to the production function, economic theory either emphasize the interaction between workers with different levels of skill, or between skills and technology. In the former case, the sorting pattern depends on whether worker skills are substitutes or complements.

If skills are complements, the marginal value of increasing the skill level of one worker is increasing in the skill level of her co-workers.⁵ For example, in Kremer (1993), one weak link –

⁵That skill complementarities can lead to positive assortative matching between workers with heterogeneous skills and firms with heterogeneous skill demands goes back at least to Becker's (1973) model of the marriage market. See also the literature on matching in labor markets with two-sided heterogeneity (Shimer and Smith, 2000 and Legros and Newman, 2002, 2007).

in the sense of a low-skilled worker – reduces the value of the production by an otherwise highly skilled chain of workers. In such a setting, a competitive labor market without search frictions ensures that workers are perfectly sorted by skill, implying that high- and low-skilled workers work in different firms.

If skills are substitutes, the marginal value of a worker’s skill is lower the more skilled are the other workers in the firm. That is, productivity hinges on the skills of a few "superstars" (Rosen, 1981) rather than a high general level of skill. In order not to waste talent, optimal sorting then implies that the most skilled workers work in different firms. Consequently, skill differences will be large within firms, and small between firms, if skills are substitutes while the converse is true if skills are complements. If skills are neither substitutes nor complements the allocation of skill across firms does not affect output, implying that sorting of workers to firms is random.⁶

The extent to which worker skills are complements or substitutes is likely to change when technology develops, although the direction of the change is not obvious a priori. In particular, it could become more important to avoid "weak links" as production processes become more complex (Kremer, 1993), suggesting that technological change increases skill complementarities. On the other hand, improvement in information technology may imply that skilled workers can leverage their skills over a wider set of problems, thereby increasing the extent to which high-skilled workers substitute for low-skilled workers.

If skills interact with technology, workers will be sorted across firms by skill to the extent that technology differ across firms. Acemoglu (1999) and Caselli (1999) develop models where skilled-biased technological change (SBTC) may shift the economy from a pooling equilibrium where firms hire both skilled and unskilled workers to a separating equilibrium where unskilled and skilled workers are sorted into different firms.⁷ In these models, SBTC thus have the same effect on sorting as an increase in the complementarity between worker skills.

Apart from changes to the production function, sorting may be affected by changes in the scope for matching workers induced by globalization. Trade in tasks, or offshoring, allows for skill-sorting by narrowing the set of tasks that needs to be performed domestically (Feenstra and Hanson, 1996; Grossman and Rossi-Hansberg, 2008). Globalization also opens up for the formation of international teams, allowing skilled workers in rich countries to match with workers in developing countries rather than unskilled workers in their own country (Kremer and Maskin 2006). The model by Grossman and Maggi (2000) opens up for a link between standard trade theory and the organization of production by letting the distribution of skills differ between countries. These differences give rise to comparative advantages in sectors where skills are either complements (supermodular) or substitutes (submodular).⁸ For a country such as Sweden, where

⁶A more formalized argument of “weak links” and “superstars” in the production function is provided in Milgrom and Roberts (1990) with the concepts of “supermodularity” and “submodularity”.

⁷There is a large literature on SBTC and its implications for the relationship between technology and skills. This literature does not, however, directly analyze worker sorting. See Acemoglu (2002), Hornsten et al (2005), and Autor and Acemoglu (2011) for surveys. Goldin and Katz (2008) provide a thorough analysis of the relation between technological change and worker skills.

⁸There is a growing theoretical literature on international trade with heterogeneous workers (e.g. Ohnsorge and Treffer, 2007; Costinot, 2009; Costinot and Vogel, 2010). These models focus on allocation between industries and not how workers with different skill levels are matched to each other.

the dispersion of skill among the workforce is low in an international comparison, the theory predicts that production of services where worker skills are complements will increase with trade, thereby increasing the optimal segregation by skill.⁹

A small empirical literature has sought to estimate whether sorting has increased over time.¹⁰ Kremer and Maskin (1996) find evidence of increased workplace segregation in France (1986-1992), the UK (1984-1990), and the US (1976-1987), using wages, experience, education, and occupation as proxies for skill. Dividing employees into production and non-production workers, Dunne et al (1997, 2004) document increases in workplace segregation in US manufacturing sectors between 1975 and 1992. Following Abowd, Kramarz, and Margolis (1999) in using worker fixed effects from a wage regression that controls for firm fixed effects as a measure of skill, Iranzo, Schivardi and Tosetti (2008) find no indication of an increase in skill sorting using data on Italian firms between 1981 and 1997. As we have argued in the introduction, there is some doubt as to whether these studies capture sorting by *skill*, rather than, for example, changes in production technology.

To the best of our knowledge, our study is the first to study sorting according to cognitive and non-cognitive skills. Our paper also differs from the previous literature in that we study a longer, and more recent, period (1986-2008). Finally, the fact that we have access to both cognitive and non-cognitive skills *and* data on years of schooling and field of study implies that we can answer questions about the underlying mechanism not possible in previous research.

3 Data

To analyze ability sorting over time, we match information on cognitive and non-cognitive skills from the Swedish military enlistment with employer-employee data. The first cohort for which we have enlistment data are men born in 1951, who were enlisted in 1969. Since it is possible to match individuals to firms in Sweden from 1986 and onwards, we can obtain a complete series of worker skill-firm matches at a given age for men at or below the age of 35. To obtain a sample of comparable individuals over time, we therefore restrict our sample in each year to men between the age of 30 and 35. We exclude men below the age of 30 from the sample to avoid a sample selection effect due to the expansion of higher education. The total sample consists of essentially all male Swedish citizens born between 1951 and 1978. Descriptive statistics for the data sets used in the analysis are available in Appendix A.

We link employees to their employers using the RAMS data base which contains information on all workers employed in a firm at some point in time each year. RAMS includes worker annual earnings by employer, the month employment started and ended, and firm level information such as ownership and industry. For workers who are recorded as having more than one employer

⁹Harrison, McLaren, and McMillan (2011) surveys the literature on trade and inequality. The literature on trade and the organization of firms is surveyed by Antràs and Rossi-Hansberg (2009). Trade and technological change are interrelated phenomena. See Bloom, Draca, and Van Reenen (2011) and references therein for theory and empirical evidence on how trade can induce technological change.

¹⁰There are also a small set of papers that study sorting in the cross-section, e.g., Hellerstein and Neumark (2008).

during a given year, we retain only the employer that corresponds to the highest annual earnings. The majority of workers only receive earnings from one firm in a given year. For example, in 2006 71 % of the workers in our data received earnings from one firm and 95 % from less than three firms. The industry classifications in RAMS have changed somewhat over time. In particular, the industry classification from 1990 onwards (SNI92) is not perfectly comparable with earlier industry classification (SNI69). We impute industry backwards 1986-1989 for firms alive in 1990. For the subsample of firms not alive in 1990, we translate 2-digit industry codes from SNI69 to SNI92.¹¹

We make some further restrictions on the sample. First, we restrict our sample to firms where we observe at least three men with complete draft records. The reason is that we are interested in studying the variation in abilities both within and between firms. Second, we restrict our sample to firms in the private sector with at least 50 employees, excluding firms controlled by the public sector and private non-profit organizations.¹² We include private firms registered in Sweden even if they are controlled from outside of Sweden, for example subsidiaries to foreign firms.

Information on basic demographics, including earnings, year of birth and educational attainment, is taken from the data base LOUISE which covers the entire Swedish population. We lack information about educational attainment prior to 1989 for about 10 percent of the sample. For this group we impute educational attainment between 1986 and 1989 using educational attainment in 1990. We translate highest educational degree into years of schooling, which we use as our measure of educational attainment (see Appendix A for details).

We obtain information on wages from the Structural Wage Statistics (SWS) which is based on annual surveys on a subsample of firms.¹³ In a given year, wages from the SWS is available for between [X] and [Y] percent of the workers in our sample. When wages are missing from the SWS, we impute wages using the SWS from other years within the same employer-employee match. For matches where no SWS wage is available, we set the wage equal to the predicted value from a regression of (observed and imputed) wages from the SWS on a high-order polynomial in the average monthly pay from RAMS. The exact details regarding our construction of wages is available in Appendix A.

3.1 Enlistment skill measures

We obtain data on cognitive and non-cognitive skills from Swedish enlistment records. The enlistment usually takes place the year a Swedish man turns 18 or 19 and spans two days involving tests of health status, physical fitness, cognitive ability, and an interview with a certified psychologist. For the cohorts we consider, the military enlistment was mandatory for all Swedish

¹¹See Appendix A for the mapping between SNI69 to SNI92.

¹²One reason for restricting attention to the private sector is that the definition of a “firm” in the public sector is not restricted to companies owned by the public sector, but also includes various types of government bodies. For example, each municipality is coded as a separate “firm” in the data.

¹³There is some variation across years in terms of the exact sampling procedure and in the number of sampled firms, but small firms are less likely to be sampled throughout our study period. More details are provided in Appendix A.

men and exemptions were only granted to men with severe physical or mental handicaps. About 90 percent of the men in our sample were eventually enlisted to the military service. Lindqvist and Vestman (2011) provide a detailed account of the enlistment procedure, the tests of cognitive ability and the enlistment interview.

Between 1969 and 1994, the enlistment test of cognitive ability consisted of four parts, testing verbal, logical, spatial and technical ability. The results of these tests were then transformed by the enlistment agency to the "stanine" scale – a discrete variable ranging from 1 to 9 that approximates a normal distribution. The basic structure of the test remained intact until 1994, although the actual test questions changed in 1980. There have also been slight changes in the mapping from raw test scores to general cognitive ability over the years (see Håkanson et al 2012 for details). A new version of the test based on the stanine scale was introduced in 1994. The youngest cohort in our main sample (men born in 1978) did the enlistment in 1996 and 1997.

We percentile-rank the 1-9 cognitive score for each set of cohorts with the same test and mapping from raw to final scores. We then convert the percentile-rank to a normally distributed test score with zero mean and unit variance. A potential concern with this procedure is that standardization hides changes in the underlying distribution of abilities. As discussed in closer detail in Appendix A, there is evidence of a "Flynn effect" – a secular rise in results on cognitive test scores over time – for logic and spatial ability while the trend for technical and verbal ability is less clear. However, except for a slight fall in the variance of verbal ability, there is no trend in the dispersion of cognitive test scores over time.


At the enlistment, conscripts were also interviewed by a certified psychologist for about 25 minutes. The objective of the interview was to assess the conscript's ability to cope with the psychological requirements of the military service and, in the extreme case, war. Each conscript was assigned a score in this respect from the same stanine scale as for cognitive ability. The instructions to the psychologists for how to evaluate conscripts was unchanged until 1995 when it was subject to slight revisions. The character traits considered beneficial by the enlistment agency include willingness to assume responsibility; independence; outgoing character; persistence; emotional stability, and power of initiative. Motivation for doing the military service was not considered beneficial for functioning in the military. We use the psychologists' evaluation as a measure of non-cognitive skill and undertake the same normalization to zero mean and unit variance as for cognitive ability.

The measures of cognitive and non-cognitive ability have a modest positive correlation (0.39), suggesting that they capture different types of ability.¹⁴ Lindqvist and Vestman (2011) show that while both skill measures predict labor market outcomes, cognitive ability is relatively more important in skilled occupations while workers in unskilled occupations have a higher return to non-cognitive ability.

¹⁴The positive correlation between cognitive and non-cognitive ability could reflect an effect of cognitive skills on non-cognitive skills, or the other way around. Lindqvist and Vestman (2011) show that the result on the cognitive test score has a small positive effect on the psychologists evaluation of conscripts' non-cognitive skills. On the other hand, noncognitive skills could affect the performance on the test of cognitive ability, as argued by Borghans, Meijers and ter Weel (2008) and Segal (2008). Moreover, noncognitive abilities could facilitate the acquisition of cognitive abilities over the life-cycle (Cunha and Heckman 2007; Cunha and Heckman 2008).

To motivate the use of the enlistment skill measures in a study of sorting, Figure 1 shows that industry wage differentials are strongly related to the average level of cognitive and non-cognitive skill. Table B1 shows that the enlistment skill measures outperform educational attainment as predictors of industry- and firm wage differentials.¹⁵

FIGURE 1: ENLISTMENT SKILL MEASURES AND INDUSTRY WAGE DIFFERENTIALS



4 Measuring sorting

We quantify sorting by decomposing the variance in skills. Let C_{ij} denote the cognitive skill of worker i in firm j . The sample variance in cognitive skill can be expressed as the sum of the variance within and between firms:

$$\underbrace{\frac{1}{N} \sum_j \sum_i (C_{ij} - C_j)^2}_{\text{within-firm variance}} + \underbrace{\frac{1}{N} \sum_j N_j (C_j - \bar{C})^2}_{\text{between-firm variance}}, \quad (1)$$

where C_j is the average level of cognitive skills in firm j , N_j is the number of workers in firm j and N is the total number of workers in the economy. In an economy where firms either hire low-skilled ("McDonald's") or high-skilled workers ("Google"), the within-firm component would be low while the between-firm component would be high. The other extreme is an economy where the average skill is the same in all firms, implying that all variance in skill is within firms. By studying the evolution of the within- and between-firm variances, we get an idea of whether sorting by skill has increased or decreased over time. Note that even though the population variances of cognitive and non-cognitive skills are set to 1 by construction, the sample variance may be either higher or lower than 1 depending on selection into the sample, i.e., private firms with at least 50 employees.

The between-firm variance can be further decomposed into variance within firms in the same industry, and variance in skill between industries. Let C_{jk} denote the average cognitive skills

¹⁵There is a long-standing debate about whether industry skill differentials reflect unobserved skill differences between workers or whether there are "true" wage differentials. See, for example, Gibbons and Katz (1992) and Gibbons, Katz, Lemieux and Parent (2005). Our objective here is not to contribute to this literature, but simply to motivate the empirical relevance of the enlistment skill measures.

of firm j in industry k and N_{jk} is the number of workers in this firm, while C_k and N_k are the corresponding variables at the industry level. The between-firm variance in cognitive skill can then be decomposed as

$$\underbrace{\frac{1}{N} \sum_k \sum_j N_{jk} (C_{jk} - C_k)^2}_{\text{variance between firms within industries}} + \underbrace{\frac{1}{N} \sum_k N_k (C_k - \bar{C})^2}_{\text{variance between industries}}. \quad (2)$$

There are a number of issues to consider regarding the use of variance decompositions as a way to measure sorting of workers to firms.

First, an implicit assumption in decomposition (1) and (2) is that we observe all workers in all firms. In fact, since we restrict attention to men between the age of 30 to 35, we observe n_j out of N_j workers in a given firm, where $n_j \leq N_j$. When $n_j < N_j$ we get a measurement error in the firm-level mean of skills, C_j , which inflates the between-firm variance and deflates the within-firm variance in (1). We show in Appendix C, that using Bessel's correction and adjusting for firm sample size implies that the decomposition in (1) becomes

$$\underbrace{\frac{1}{n} \sum_j n_j \left(\frac{N_j - 1}{N_j} \right) \left(\frac{1}{n_j - 1} \right) \sum_i (C_{ij} - C_j)^2}_{\text{within-firm variance}} + \underbrace{\frac{1}{n} \sum_j n_j \left[(C_j - \bar{C})^2 - \frac{N_j - n_j}{N_j n_j} \left(\frac{1}{n_j - 1} \right) \sum_i (C_{ij} - C_j)^2 \right]}_{\text{between-firm variance}}. \quad (1')$$

All results presented in the paper are based on decompositions that adjust for sample size, but, to save on space, we show the expressions in Appendix C.¹⁶ Note that we have chosen to weigh each firm by the number of observed workers (n_j) rather than the actual number (N_j) in (1').¹⁷

Second, since the number of workers at each firms is finite, the between-firm variance would be larger than zero also under random matching. To get a benchmark value of sorting, we randomly draw workers to firms without replacement from the set of workers in the sample and conduct the variance decomposition in (1). Repeating this process 1,000 times provides a bootstrap-type test

¹⁶The adjustment in (1) implies that we view each firm as a population rather than a sample of random draws from a given distribution of worker skills. It is not obvious a priori which view is most accurate as sorting arguably both reflect random and deterministic factors. A potential concern with viewing the set of workers in a firms as a population rather than a sample is that a shift in the size distribution of firms could lead to changes in the estimated within- and between-firm components even if worker skills are drawn from the same underlying distribution. However, as discussed in the next paragraph, we simulate benchmark values of each component in that takes shifts in the size distribution of firms into account.

¹⁷There are two reasons for this choice. First, weighting firms by the number of observed workers is more efficient. Weighting firms by the actual number of workers would imply that a number of firms with few observed workers would get a large weight, thus increasing random noise. Second, since our sample is restricted to men in the age of 30-35 in the first place, weighting firms by the actual number of workers would not be representative of the entire population of workers unless one is willing to assume that sorting patterns are exactly identical for 30-35 year old men compared to the population as a whole.

of sorting by comparing the true between-firm variance with the percentiles in the distribution of simulated variances.¹⁸ Comparing the actual and simulated between-firm variances is a first simple test of what forces drive sorting in the aggregate. If worker skills are complements, or if there is a complementarity between worker skills and technology (and technology differs across firms), then the actual between-firm variance should exceed the simulated variances. If, in contrast, there are no or weak complementarities between worker skills and technology and worker skills are substitutes, the observed level of sorting should be below the simulated level.¹⁹

Third, the enlistment skill measures are likely affected by measurement error. Using data on monozygotic and dizygotic twins, Lindqvist and Vestman (2011) estimate a reliability ratio of 0.868 for cognitive and 0.703 for non-cognitive skills.²⁰ As shown in Appendix C, measurement error inflates the within-firm variance relative to the between-firm variance. Since the effect of measurement error on the estimated firm mean of skills is smaller the larger are firms, a change in the size distribution of firms over time could affect the share of the measurement error variance that is attributed to within- and between-firm components. Assuming classical measurement error, we derive a correction for measurement error. In essence, we use the estimated reliability ratios from Lindqvist and Vestman (2011) to simulate measurement errors for each worker in our data. We then use the simulated errors to estimate the share of the within- and between-firm variance which can be attributed to measurement error (see Appendix C for details). We report these results as a robustness check rather than as our main case.

Fourth, we assume that the enlistment skill measures follow a normal distribution. Although a reasonable benchmark case, there is no fundamental argument as to why skills should be normally distributed. It is thus fair to ask how robust our results are to monotone transformations of skills or non-parametric ways of quantifying sorting. To test the sensitivity to distributional assumptions, we transform the enlistment skill measures alternative distributions (uniform and Beta distributions with different skew) which we then decompose into between- and within firm components. To estimate sorting non-parametrically, we first rank all firms in each year according to the average level of skills. We then calculate the Kendall's tau rank correlation between the rank of each firm and skill level of each individual. As with the standard variance decomposition, we compare the results for non-parametric sorting with randomly generated data sets.²¹

Finally, our sample is restricted to men between 30 and 35. An advantage of this restriction is that the high mobility of young male workers is that we are likely to detect changes in sorting patterns quickly. Still, the external validity would be stronger if the same sorting patterns are present for female workers and older male workers. We impute cognitive and non-cognitive

¹⁸A similar approach is used by Ahlin (2010).

¹⁹Since the counterfactual is simulated on the sample, the within-firm variance is just the residual of the between-firm variance. The choice to draw workers from the sample rather than the population is not obvious since the sample is arguably not exogenous. However, since the sample variance of skill is lower than the population variance, drawing from the sample biases our results against positive sorting. Relatedly, we draw workers from the aggregate (non-stratified) sample, implying that a worker may be matched to a firm in an industry in which the worker does not actually work.

²⁰The lower reliability ratio of non-cognitive skills arguably reflects the additional error introduced by the fact that different psychologists evaluate different conscripts (Lindqvist and Vestman, 2011).

²¹Our approach for quantifying sorting using Kendall's tau is similar to Ahlin (2010).

skills for women using the draft records of close male relatives (see Appendix A for details). We then decompose the variance in skills for women between the age of 30 and 35 following the same procedure as for men. Since we have to impute cognitive and non-cognitive skills of females, measurement error in skill is an order of magnitude larger for this sample, leading to a spuriously low level of sorting across firms. We discuss our approach for adjusting for measurement error in the female sample in Appendix C. To test the robustness of our results with respect to age, we study the sorting patterns from 1996 to 2008 for male workers between the age of 30 and 45.

The variance decompositions in (1) and (2) quantify sorting in each skill measure separately. A related question is how the covariance between the firm-level averages of different types of skill has evolved over time. Just like the variance, the covariance of cognitive and non-cognitive skills can be decomposed into between- and within-firm components. Let C_{ij} and NCS_{ij} denote cognitive and non-cognitive skills of worker i in firm j and C_j and NCS_j the corresponding averages for firm j . We can then decompose the covariance between cognitive and non-cognitive skill as

$$\underbrace{\frac{1}{N} \sum_j \sum_i (C_{ij} - C_j) (NCS_{ij} - NCS_j)}_{\text{within-firm covariance}} + \underbrace{\frac{1}{N} \sum_j N_j (C_j - \bar{C}) (NCS_j - \overline{NCS})}_{\text{between-firm covariance}}.^{22} \quad (3)$$

The between-firm covariance tells us to what extent firms whose workers have high cognitive skills also have high non-cognitive skills. Since cognitive and non-cognitive skills are positively correlated at the level of the individual, the sum of the within- and between-firm components is positive. However, depending on how skills are valued across firms, the between-firm covariance could in principle be negative. For example, if the return to skill are negatively correlated across sectors, firms in sectors that value cognitive skill will hire workers who have relatively low non-cognitive skills given their cognitive skills, since these workers command a lower price in the market for labor. The same sorting pattern would occur if cognitive and non-cognitive skills are substitutes in the firm-level production function.²³ Still, the fact that cognitive and non-cognitive skills are positively correlated at the individual level does imply that the between firm covariance will be positive in our benchmark case where workers are sorted randomly to firms.

5 Sorting by skill 1986-2008

In this section, we document the evolution of skill sorting in the Swedish economy over the last 25 years. We begin with the most basic question: Has skill sorting increased or decreased?


Figure 2 shows the evolution of the within- and between-firm variance for the enlistment skill

²²The decomposition that adjusts for sample size in firms is presented in Appendix E.

²³In a future work, we plan to estimate production functions augmented with cognitive and non-cognitive skills, thereby allowing us to infer whether cognitive and non-cognitive skills are negative or positive complements, and whether returns to skills are positively or negatively correlated across different sectors of the economy. An alternative approach would be to estimate the production function that fits best with the observed sorting pattern along the lines suggested by Fox (2010). However, both of these exercises are beyond the scope of the present paper.

measures between 1986 and 2008. Panel A shows that the within-firm variance in cognitive skill fell from 0.80 in 1986 to 0.70 in 2008. At the same time, the between-firm variance increased from 0.14 to 0.19. We can thus conclude that sorting has increased: people working in the same firm have become more similar while workers in different firms have grown more different in terms of their cognitive skills. The reason the fall in the within-firm variance is not fully reflected in a corresponding increase in the between-firm variance is that the variance of cognitive skills in our sample falls somewhat during the study period. As shown in Panel B, the trend for non-cognitive skills is similar to that of cognitive skills, even though the between-firm variance is lower. As it turns out, all of the simulated variances are well below the estimated variances for all skill measures, suggesting that complementarities dominate any "superstar" effect.

FIGURE 2: BETWEEN- AND WITHIN-FIRM VARIANCE IN SKILL 1986-2008



Note: Sample men 30-35 years old, firms with at least 50 employees

[INCLUDE ROBUSTNESS TESTS HERE]

In the next section, we proceed to take a closer look at each part of the decomposed variance starting with the within-firm variance.

5.1 Decomposing the within firm variance

To get a more detailed picture of the driving forces behind the fall in within-firm variance, we decompose the change in within-firm variance as

$$\sum_k \alpha_{k,86} \Delta \sigma_k^2 + \sum_k \Delta \alpha_k \sigma_{k,86}^2 + \sum_k \Delta \alpha_k \Delta \sigma_k^2, \quad (4)$$

where $\alpha_{k,86} = N_{k,86}/N_{86}$ denotes the share of the sample employed in industry k in 1986, $\sigma_{k,86}^2$

is the average within-firm variance (weighted by firm size) for the same year and industry, $\Delta\sigma_k^2 = \sigma_{k,08}^2 - \sigma_{k,86}^2$ and $\Delta\alpha_k = \alpha_{k,08} - \alpha_{k,86}$. The first term in (4) is the change in within-firm variance holding each industry's share of total employment fixed at its 1986 level. This term should be negative if increasing complementarities between skills in the production function or diffusion of new technology makes it more profitable to match workers of a given skill level in the same firm. The second term is the change in within-firm variance which is due to a change in the relative size of industries. If, as suggested by Grossman and Maggi (2000), Sweden has a comparative advantage in goods and services where worker skills are complements, falling trade costs should lead to an increase in the relative size of industries where the initial within-firm variance ($\sigma_{k,86}^2$) is small and, consequently, a negative second term. The third term is the covariance between changes in the relative size of industries and changes in within-firm variance.

Table 1 displays decomposition (4) for each of our skill measures. The fall in the within-firm variance is mostly due to a fall in the within-firm variance for fixed industry shares. Industries with a low initial within-firm variance increased in size relative to other industries, but this effect can only explain a small share of the overall trend. The results thus do not support lower trade costs as the main driving force behind the fall in the within-firm variance of skill.

TABLE 1. DECOMPOSING CHANGE IN WITHIN-FIRM VARIANCE

Skill measure	Δ Within-firm variance	Δ Size of industries	Covariance
Cognitive	-0.088	-0.030	0.012
Non-cognitive	-0.058	0.002	-0.007

In which industries did the within firm variance in skill fall the most? Table 2 lists the within-firm variance (weighted by firm size) by industry for each skill measure, the change in within-firm variance between 1986 and 2008, the relative size of the industry in 1986 and the change in relative size. We restrict the sample to industries with at least 2 percent of the workforce in either 1986 or 2008. The industries in the table are sorted according to average within-firm variance (weighted by firm size) in 1986. Table 2 shows that the within-firm variance in cognitive skill fell sharply in a range of manufacturing industries, including manufacture of telecom products; the chemical industry; pulp and paper; and the forest industry. However, the fall in within-firm variance is a feature of almost all industries and holds for all skill measures. The main exception is the telecommunications industry, but the diminutive size of this industry in 1986 implies that one should be cautious in interpreting this result. Table 2 also shows that the fall in within-firm variance is partly caused by the growth of the IT sector ("Computer and related activities"), which has the lowest within-firm variance of cognitive skill among the major industries.

The general trend toward smaller within firm variance is present also for non-cognitive skills, albeit not as dramatic as for cognitive skills. Unlike cognitive skills, large shifts in the within-firm variance of non-cognitive skills does not only pertain to manufacturing industries.

TABLE 2. WITHIN-FIRM VARIANCE IN SKILLS 1986 AND 2008

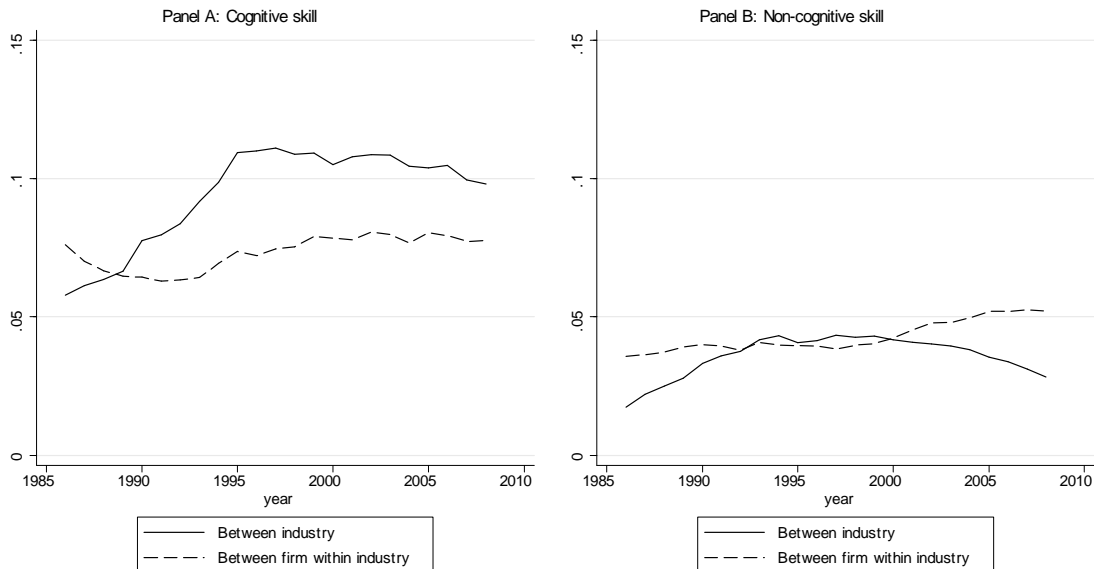
Industry	Cognitive		Non-cognitive		α_{1986}	$\Delta\alpha$
	$\sigma_{k,86}^2$	$\Delta\sigma_k^2$	$\sigma_{k,86}^2$	$\Delta\sigma_k^2$		
Manufacture of motor vehicles, trailers and semitrailers	0.97	-0.11	0.84	-0.08	6.71	-0.35
Manufacture of chemicals and chemical products	0.96	-0.15	0.88	-0.15	2.13	-0.14
Manufacture of radio, television and communication equipment	0.93	-0.25	0.82	-0.14	2.57	0.08
Manufacture of food products and beverages	0.89	-0.09	0.84	-0.06	4.03	-1.39
Manufacture of machinery and equipment n.e.c	0.88	-0.08	0.76	0.00	7.96	-1.16
Manufacture of wood and of products of wood and cork	0.86	-0.13	0.75	-0.07	2.27	-0.71
Manufacture of basic metals	0.85	-0.08	0.76	0.07	2.39	-0.21
Manufacture of fabricated metal	0.85	-0.10	0.79	-0.05	3.68	-1.11
Manufacture of paper and paper products	0.85	-0.14	0.82	-0.03	5.36	-3.61
Publishing, printing and reproduction of recorded media	0.79	-0.09	0.96	-0.08	2.74	-1.38
Other business activities	0.77	-0.11	0.81	-0.05	8.40	4.12
Land transport; transport via pipeline	0.76	0.02	0.85	-0.17	1.43	0.67
Wholesale trade	0.75	-0.11	0.82	-0.07	9.46	-2.44
Retail trade, except of motor vehicles	0.72	-0.01	0.82	-0.01	2.11	2.07
Supporting and auxiliary transportation activities	0.72	0.01	0.88	-0.11	1.67	0.59
Construction	0.70	-0.09	0.76	-0.07	8.82	-0.14
Sale, maintenance and repair of motor vehicles	0.66	-0.03	0.73	-0.01	2.65	0.22
Financial intermediation, except	0.59	-0.03	0.81	-0.08	3.48	-0.74
Computer and related activities	0.59	-0.02	0.75	-0.01	1.20	6.66
Post and telecommunications	0.49	0.24	0.84	-0.05	0.02	2.15

Sample restricted to firms with at least 2 % of the workforce in either 1986 or 2008.

5.2 Decomposing the between firm variance

We now turn to a more in-depth analysis of the between-firm variance of skill. Figure 3 decomposes the between-firm variance in skills into skill differences between industries, and differences in average skill between firms within the same industry. We document a substantial increase, from 0.06 to 0.11, in the between-industry variance of cognitive skill from 1986 to 1995. The pattern is similar for non-cognitive skills up until the mid 1990's when the between-industry variance fell somewhat. In general, sorting at the industry level appears to be more important for cognitive than for non-cognitive skills. This result is consistent with the finding in Lindqvist and Vestman (2011) that cognitive skills is a stronger predictor of selection into skilled or unskilled occupations than non-cognitive skills. Figure 3 also shows that the variance in skill between firms within the same industry increases from the mid 1990's and onwards.

FIGURE 3: DECOMPOSING THE BETWEEN-FIRM VARIANCE 1986-2008



Note: Sample men 30-35 years old, firms with at least 50 employees

The fact that sorting at a relatively coarse industry level (about 50 industries) explains more than half of the between-firm variance in cognitive skill and education suggests that broad differences in technology are important for explaining sorting by skill, and also have become more important over time. To get a more detailed picture, we decompose changes in the between-industry variance as

$$\sum_k \alpha_{k,86} \Delta (C_k^{cen})^2 + \sum_k \Delta \alpha (C_{k,86}^{cen})^2 + \sum_k \Delta \alpha_k \Delta (C_k^{cen})^2 \quad (5)$$

where $\alpha_{k,86} = N_{k,86}/N_{86}$ denotes the share of the sample employed in industry k in 1986, $C_{k,86}^{cen} = C_{k,86} - \bar{C}_{86}$ is the average cognitive skill in industry k centered by the population average, $\Delta (C_k^{cen})^2 = (C_{k,08}^{cen})^2 - (C_{k,86}^{cen})^2$ and $\Delta \alpha_k = \alpha_{k,08} - \alpha_{k,86}$. Table 3 shows the results from

decomposition (5). The bulk of the increase in between-firm variance of skill is due to an increase in the employment share of industries with values of skills away from the sample mean, while a smaller part is attributed to shifts in the mean values of skill away from the mean. The negative covariance term implies that industries which grew in size moved toward the sample mean of skills.

TABLE 3. DECOMPOSING CHANGE IN BETWEEN-FIRM VARIANCE

Skill measure	Δ Industry means	Δ Size of industries	Covariance
Cognitive	.015	.033	-.009
Non-cognitive	.008	.015	-.014

Which industries grew in size? Table 4 lists the mean values by industry for each skill measure, the change in mean between 1986 and 2008 and employment shares for the same set of major industries as in Table 2.²⁴ The fact that stand out from Table 4 is the rapid growth of the IT industry, the industry with the highest level of cognitive skill in both 1986 and 2008. In 2008, 7.8 % of 30-35 year old men worked in this industry, up from 1.2 % in 1986. Despite the increase in size, the average cognitive ability of workers in the IT sector remained constant at 0.7 standard deviations above average. The upshot is that the growth of the IT industry explains about two thirds of the overall increase in the between-industry variance of cognitive skill.²⁵ The remaining third is, for the most part, explained by a decline of cognitive skill in a number of low-skilled service industries, including retail, construction, transportation, sales and repair of motor vehicles. In sum, the increase in sorting is due to an increase in the number of high-skilled firms ("Google") and a fall in the skill level of the workers at low-skilled firms ("McDonald's").²⁶ Figure 4 shows the distribution of industry mean values of cognitive skill in 1986 and 2008 with each industry weighted by its employment share. The increase in density at the high end of cognitive skill brought about by the increase in the IT industry implies that the


²⁴The industry with the highest average level of cognitive skills - research and development - is not included in the Table 4 as it employes less than 2 percent of the workforce.

²⁵This result holds regardless of whether we calculate counterfactual between-industry variances excluding the IT industry keeping the sample mean of cognitive skills fixed, or removing the IT sector altogether. **[Check]**

²⁶Telecom services explain about half of the negative covariance between changes in size and the square of centered skills. This industry was tiny in 1986 (0.02 %) but had a high average skill level. Since the manifold expansion in size up to 2008 was accompanied by a fall in average skill, the covariance term is strongly negative for this particular industry.

distribution is substantially more polarized in 2008.

FIGURE 4. MEAN COGNITIVE SKILL AT THE INDUSTRY LEVEL



Note: Sample men 30-35 years old, firms with at least 50 employees

The growth of the IT industry and fall in the skill level in low-skilled industries also contribute to the increase in sorting of workers with respect to non-cognitive skill. However, the between-industry variance in non-cognitive skill is also caused by a substantial upgrading of skill in financial intermediation.

5.3 Covariance in skills

We now turn to the covariance between cognitive and non-cognitive ability. Figure 5 shows the decomposed covariance over time for services and manufacturing industries separately. The between-firm covariance is positive in both sectors, but significantly higher in services. This could reflect a stronger complementarity between skills in services compared to manufacturing, or a stronger correlation in the returns to cognitive and non-cognitive skill across firms in the services sector. Figure 5 also shows that the between-firm covariance increased (while the within-firm covariance fell) from the mid 1980's up to the mid 1990's.

The fact that the between-firm covariance of cognitive and non-cognitive skill is positive and increasing over our study period has important implications for the overall picture of sorting. In principle, an increase in between-firm differences in cognitive and non-cognitive skill could occur alongside a fall in the between-firm covariance. A fall in the between firm covariance would suggest that firms to an increasing extent hire workers with a particular type of skill, rather than

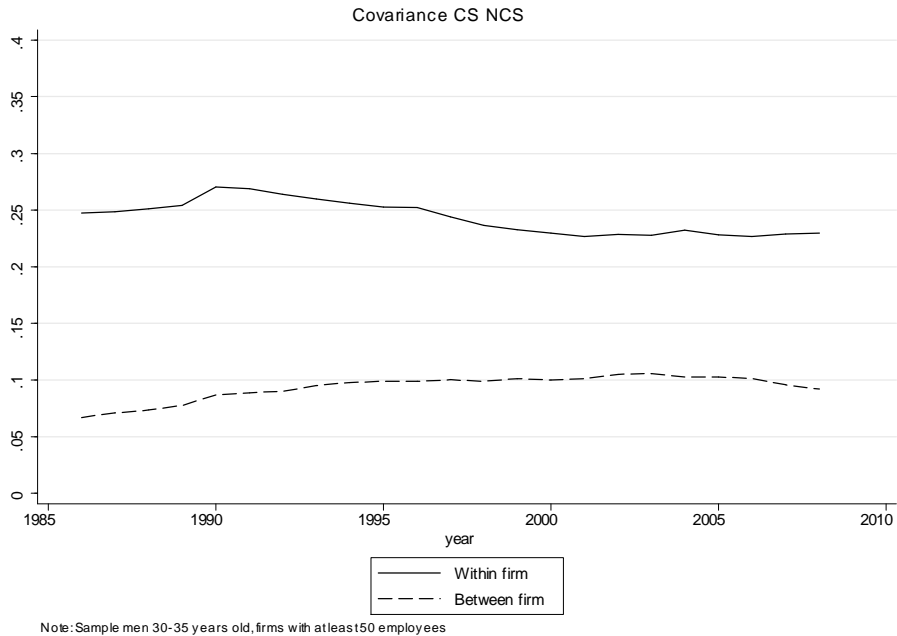
TABLE 4. MEAN VALUES OF SKILL BY INDUSTRY 1986 AND 2008

	CS_{1986}	ΔCS	NCS_{1986}	ΔNCS	α_{1986}	$\Delta\alpha$
Computer and related activities	0.70	0.00	0.23	0.04	1.20	6.62
Post and telecommunications	0.56	-0.31	0.76	-0.65	0.02	2.15
Manufacture of radio, television and communication equip.	0.46	0.14	0.09	0.14	2.55	0.09
Financial intermediation, except insurance and pension funding	0.30	0.10	0.22	0.24	3.49	-0.62
Other business activities	0.22	0.04	0.10	0.04	8.41	4.10
Wholesale trade and commission	0.15	-0.21	0.12	-0.05	9.48	-2.48
Publishing, printing and reproduction of recorded media	0.12	0.04	-0.06	0.01	2.74	-1.39
Manufacture of chemicals and chemical products	0.11	0.05	0.01	0.12	2.13	-0.13
Supporting and auxiliary transport	0.08	-0.25	0.07	-0.14	1.66	0.58
Manufacture of motor vehicles	-0.03	-0.03	-0.07	-0.00	6.67	-0.26
Manufacture of machinery and equipment	-0.06	0.07	-0.06	0.06	7.92	-1.16
Retail trade, except of motor vehicles	-0.08	-0.06	-0.06	-0.00	2.10	2.06
Sale, maintenance and repair of motor vehicles	-0.13	-0.15	-0.04	-0.11	2.64	0.21
Construction	-0.16	-0.03	0.01	0.03	8.89	-0.27
Manufacture of fabricated metal	-0.19	-0.05	-0.17	0.01	3.67	-1.12
Manufacture of paper and paper products	-0.20	0.06	-0.08	0.07	5.35	-3.61
Manufacture of food products and beverages	-0.22	-0.13	-0.15	-0.01	4.01	-1.35
Land transport; transport via pipelines	-0.27	-0.13	-0.28	-0.06	1.43	0.65
Manufacture of basic metals	-0.32	0.07	-0.18	0.07	2.38	-0.22
Manufacture of wood and of products	-0.37	-0.00	-0.16	0.03	2.26	-0.71

Sample restricted to firms with at least 2 % of the workforce in either 1986 or 2008.

workers with a high general skill level. It is not possible to tell these different stories apart with access to only a single measure of skill.

FIGURE 5. COVARIANCE BETWEEN COGNITIVE AND NON-COGNITIVE SKILL



Before we continue with a more in-depth analysis of the sorting patterns documented in the previous section, we discuss sorting according to educational attainment.

5.4 Educational attainment

Unlike the enlistment skill measures, educational attainment has a natural metric in years of schooling. However, as we have argued above, a potential problem with education as a measure of skill is that the expansion of higher education implies that an absolute measure of educational attainment (such as years of schooling) is not comparable over time. We therefore also construct an alternative measure where we first percentile-rank all individuals within each cohort with respect to their years of schooling. As for the enlistment skill measures, we then convert the percentile-rank to a normally distributed score with zero mean and unit variance.

Figure 6 shows the changes in the between and within-firm variance using years of schooling and our normalized measure of skill. Panel A shows that the within-firm variance for years of schooling first fell significantly but then increased from year 2000 and onwards. The between-firm variance in years of schooling increased throughout our study period, but the timing is different compared to cognitive and non-cognitive skill.

[INCLUDE FIGURE 6 HERE]

6 Mechanisms

So far, we documented two basic facts about changes in sorting over time in the Swedish labor market. First, across all industries, firms have become more homogeneous with respect to the skills of their workforces. Second, following the rise of the IT industry, skill differences between industries have increased.

In this section, we provide suggestive evidence regarding the mechanisms at play. In particular, we decompose the between- and within-firm components further in order to separate between skill differences due to technology, assortative matching of worker skills for a given technology, and the interaction between the two. Our starting point is that information on occupational structure can be used to construct a measure of the skill-intensity of technology. Using this measure, we then test whether the increase in the between-firm variance of skill is explained by increasing differences in the skill-intensity of technology across firms, or stronger assortative matching of workers. Thereafter, we perform the corresponding analysis for the within-firm variance.

6.1 Between-firm variance

Let C_h denote the average cognitive skills of workers in occupation h and N_{hj} the number of workers in occupation h in firm j . Then $\widehat{C}_j = \frac{1}{N_j} \sum_h N_{hj} C_h$ is the expected firm-level mean of cognitive skills in firm j conditional on the firm's occupational structure. We argue that \widehat{C}_j can be thought of as a proxy for the skill-intensity of technology in a firm. For example, a firm that hires many engineers, who are high-skilled on average, will have a high value of \widehat{C}_j . Note that two firms could have the same value of \widehat{C}_j while still producing different goods or services. For example, a law firm and a computer consultancy could be expected to have the same average skill level.

Since occupation data is missing for blue-collar workers between 1986 and 1995, we use the combination of field of study and years of schooling as a proxy for occupation. For example, workers with a 5-year degree in engineering will be coded as belonging to the same "occupation". Even though data on occupation for all sectors is available from 1996, we choose to stick to our proxy for consistency. In practice, there it does not seem to matter much whether we calculate \widehat{C}_j based upon occupation or education: The correlation is over 0.90 in 1996 and, as shown in Appendix B, the results come out in a similar fashion for the 1996-2008 period when we use \widehat{C}_j based upon occupation.

Using \widehat{C}_j , we decompose the between-firm variance of cognitive skill as²⁷

²⁷Another way of expressing (7) would be to replace $C_j - \widehat{C}_j$ with $\widehat{\varepsilon}_j = \frac{1}{N_j} \sum_i \widehat{\varepsilon}_{ij}$, where $\widehat{\varepsilon}_{ij}$ is the residual of worker i in firm j from regressing C_{ij} on occupational fixed effects. Similarly, \widehat{C}_j is the firm-average of the predicted values from this regression.

$$\frac{1}{N} \sum_j N_j \left(\underbrace{(C_j - \hat{C}_j)^2}_{\text{assortative matching}} + \underbrace{(\hat{C}_j - \bar{C})^2}_{\text{skill-intensity}} + 2 \underbrace{(C_j - \hat{C}_j)(\hat{C}_j - \bar{C})}_{\text{skill-tech complementarity}} \right). \quad (6)$$

The first term in (6) is the between-firm variance in the difference between the actual and predicted mean of cognitive skills. This term reflects assortative matching of workers within and between occupational groups in the same firm. For example, assortative matching is positive if the most and least clever engineers work in different firms, and if the most clever engineers work with the most clever secretaries. If skill complementarities between workers become stronger, we expect assortative matching to play a larger role.

The second term in (6) is the variance in skill attributable to differences in the skill-intensity of technology across firms. An increase in this terms could reflect skilled-biased technological change. For example, in the model by Caselli (1999), all workers work in the same type of firm and use the same technology in the pooling equilibrium. Skilled-biased technological change leads to a separating equilibrium where workers sort by skill to different firms and work with different technology. The second term could also increase due to outsourcing. For example, consider a firm which both develops new products (skill-intensive) and manufactures them (not skill intensive). If product development and manufacturing is instead split into two different firms, the result would be an increase in the between-firm variance due to differences in the skill-intensity of technology.

The third term is the covariance between the first and second terms. The covariance is positive if firms that employ workers from qualified occupations also employ the most skilled workers within each occupation. For example, the covariance could be positive if firms that employ many engineers (who are high-skilled on average) also employ the most skilled engineers and the most skilled secretaries. The covariance could be thought of as a measure of the complementarity between worker skills and technology.

FIGURE 7: DECOMPOSING THE BETWEEN-FIRM VARIANCE

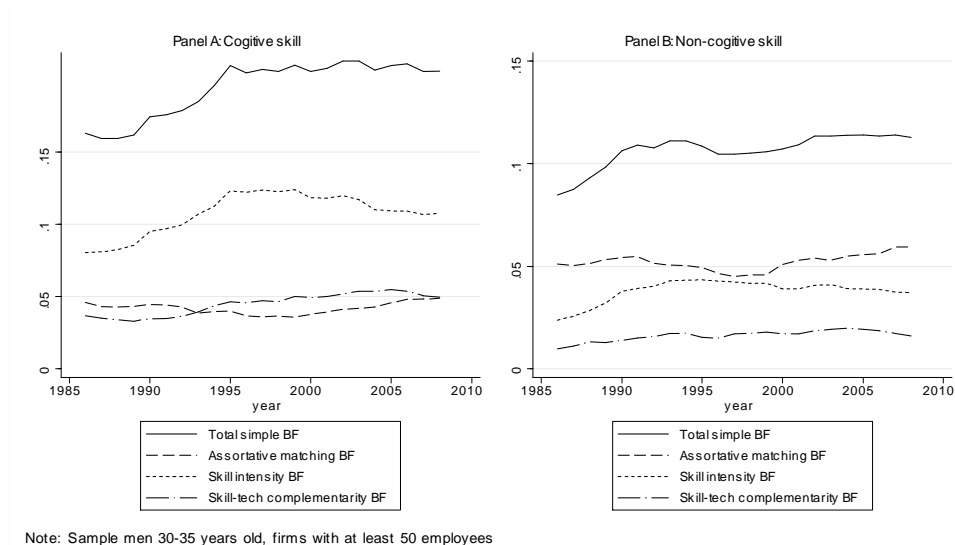


Figure 7 shows the change of each component in (6) between 1986 and 2008. Clearly, the bulk of the increase in the between-firm variance of cognitive skill between 1986 and 1995 is explained by larger differences in the skill-intensity of technology across firms. The level of assortative matching in cognitive skill is larger than predicted by random sorting, indicating that cognitive skills are complements in the production function also for a given technology, but does not change much over time. The skill-technology complementarity is positive – suggesting that firms with a skill-intensive technology level also hire the best workers in each occupation – and somewhat increasing over time. All three estimated components in (6) are "statistically significant" in the sense that they are extremely unlikely to be generated by random sorting.²⁸ The time-pattern for non-cognitive skill is similar to cognitive skill, but there is a notable level difference in the sense that assortative matching of workers is more important than the skill-intensity of technology for explaining between-firm differences in skill.

Why did differences across firms in the skill-intensity of technology increase? The finding from Section 5 that the rise of the IT industry explains a large fraction of the rise in the between-firm variance of cognitive suggests that technological change rather than outsourcing drive this result. In order to get at the underlying mechanism, we compute counterfactual between-firm variance in the skill-intensity of technology removing different groups of workers from the data.²⁹ Figure 8 shows the results from four such counterfactual decompositions.

[INCLUDE FIGURE 8 HERE]

Combined with the results from Section 2, the results above suggest a simple story for the increase in between-firm differences in cognitive skill: Following the arrival of a new technology

²⁸We use the residuals and predicted values from a regression of skills on occupation proxies as the basis for this simulation.

²⁹In the results reported in Figure 10, the sample means used to calculate the variance does not include the groups removed from the data. [Robustness check - think more about what is appropriate here]

(IT), high-skilled workers (in particular engineers and technicians) sort into the IT industry.

[DISCUSS RESULTS FOR NON-COGNITIVE SKILLS]

6.2 Within-firm variance

We now turn to an in-depth analysis of the evolution of the within-firm variance. Expression (7) shows the decomposition of the within-firm variance of cognitive skill. Let \widehat{C}_i denote the average cognitive skill in the occupation held by individual i while C_i denotes worker i 's actual skill level. Consequently, $C_i - \widehat{C}_i$ equals worker i 's residual of a regression of actual skills on occupation fixed effects. We get

$$\frac{1}{N} \sum_j \sum_i \left[\underbrace{\left(\widehat{C}_i - \widehat{C}_j \right)^2}_{\text{skill-intensity}} + \underbrace{\left(\left(C_i - \widehat{C}_i \right) - \left(C_j - \widehat{C}_j \right) \right)^2}_{\text{assortative matching}} \right] + \underbrace{2 \left(\widehat{C}_i - \widehat{C}_j \right) \left(\left(C_i - \widehat{C}_i \right) - \left(C_j - \widehat{C}_j \right) \right)}_{\text{skill-tech complementarity}}. \quad (7)$$

The first term of (7) captures the within-firm variance in the skill intensity of jobs. A fall in this term implies that firms concentrate on a set of tasks with similar skill requirements. In particular, a trend toward outsourcing and focus on core activities should lead to a fall in within-firm differences in the skill requirements of jobs.

The second term of (7) is large if there is a high variance of worker skills within occupations, or between occupational groups compared to the relative skill level of the firm. For example, the second term would be high for a firm where all engineers were more skilled compared to engineers in general whereas production workers were unskilled compared to the average production worker. The second term is smaller the stronger is positive assortative matching.

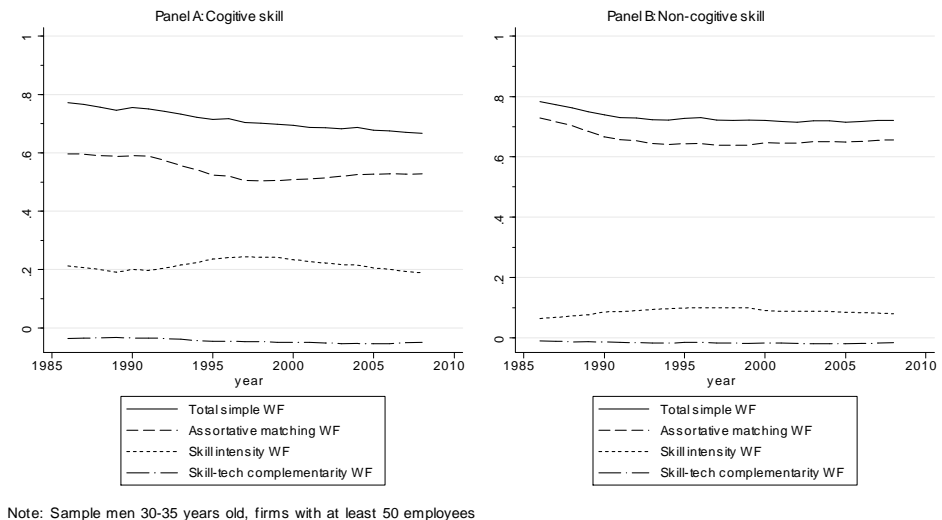
The last term in this expression is the mirror image of the covariance term in (6) in the sense that it has the same absolute value but the opposite sign.³⁰ Since the covariance term in (6) is positive, it follows that the covariance term in (7) is negative. In words, if firms that hire a larger proportion of workers from high-skilled occupations (e.g., engineers) also hire the most skilled workers in each occupation, then *within firms*, workers in skilled occupations will be less skilled *relative* to their occupations.

Figure 9 shows that the fall in within-firm variance of cognitive and non-cognitive skill is mainly due to stronger positive assortative matching. In contrast, there is no trend toward smaller variance in skill requirements within firms. In other words, workers in the same firm have not become more similar because they the range of technology used within the firm has narrowed. We view this as suggestive evidence against outsourcing since, if firms had outsourced part of their business, we would expect the skill requirements of workers in the same firm to

³⁰We show that the third term in (7) equals the third term in (8) multiplied by -1 in Appendix C.

become more homogeneous.

FIGURE 9 : DECOMPOSING THE WITHIN-FIRM VARIANCE



Why has assortative matching increased? In order to investigate this issue, we regress our measure of assortative matching component on different sets of covariates. Specifically, let $\sigma_{Ass,jkt}^2$ denote the second component from decomposition (7) for firm j in industry k at time t . We estimate regressions of the following generic form

$$\sigma_{Ass,jkt}^2 = \beta_0 + \beta_1 \log(Capital)_{jkt} + \beta_2 \log(Size)_{jkt} + \beta_3 (C_{jkt} - \hat{C}_{jkt}) + \beta_4 \hat{C}_{jkt} \quad (8) \\ + \beta_5 Manufacturing_{kt} + \beta_6 Trade_{kt} + \beta_7 China_import_{kt} + \varepsilon_{jkt}$$

where $Capital_j$ is capital intensity, $Size_j$ is the number of employees, $(C_{jkt} - \hat{C}_{jkt})$ is the difference between the actual and predicted skill level of firm j and \hat{C}_{jkt} is the predicted skill level of firm j based upon its capital structure and $Manufacturing_{kt}$ an indicator for whether industry k belongs to the manufacturing sector. For a subset of industries (mainly in manufacturing), we also have rough data on trade. $Trade_{kt}$ equals the total value of exports and imports in industry k divided by total value added while $China_import_{kt}$ equals imports from China divided by value added. We think of $China_import_{kt}$ as a proxy for competition from low-wage countries, and also the scope for outsourcing production to other countries (Kremer and Maskin, 2006).

In regression (8), β_1 and β_5 reflects the association between the type of production process and assortative matching, controlling for the skill-intensity of technology. β_4 answers the question whether assortative matching is correlated with the skill-intensity of technology. If more complex production processes are associated with stronger complementarities (as in Kremer 1993), then we expect $\beta_4 < 0$. β_3 answers the question whether "star" firms with unexpectedly high skills also have stronger assortative matching. Note that $\beta_3 + \beta_4$ gives the total "effect" of C_j on assortative matching.

We estimate regression (8) using both cross-sections at the beginning (1986) and end (2008) of our study period, as well as fixed-effects estimation using the panel for the entire 1986-2008 period. We also collapse the values of all variables in (8) at the industry-level and estimate the long difference in assortative matching on the long differences (and levels in 1986) of the right-hand side variables in (8). Since we only have trade data for a subset of industries (mainly in the manufacturing sector), we estimate all versions of (8) using both this subset and the whole sample excluding the trade variables. We weight each firm or industry is weighted with the number of workers observed in our sample.

[INCLUDE RESULTS HERE]

7 Conclusions

[TO BE WRITTEN]

8 References

Abowd, J. M., F. Kramarz and D. N. Margolis (1999), "High Wage Workers and High Wage Firms," *Econometrica*, 67(2), 251-333.

Acemoglu, D. "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence," *American Economic Review*, 1999, 89(5), 1259-1278.

Acemoglu, D. "Technical Change, Inequality, and the Labor Market," *Journal of Economic Literature*, 2002, 40(1), 7-72.

Antràs, P. and E. Rossi-Hansberg (2009), "Organizations and Trade", *Annual Review of Economics*, 1, 43-64.

Autor and Acemoglu (2011), "Skills, Tasks and Technologies: Implications for Employment and Earnings," *Handbook of Labor Economics 4B*, 2011, 1043-1171.

Becker, G. S. (1973), "A Theory of Marriage: Part I," *Journal of Political Economy*, 81(4), 813-846.

Bloom, N., M. Draca, and J. Van Reenen (2011), "Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT, and Productivity," CEP Discussion Paper No 1000.

Borghans, L., B. ter Weel and B. A. Weinberg. "Interpersonal Styles and Labor Market Outcomes," *Journal of Human Resources*, 2008a, 53(4), 815-858.

Borghans, L., A. L. Duckworth, J. J. Heckman, and Bas ter Weel. 2008b. "The Economics and Psychology of Personality Traits." *Journal of Human Resources*, 53(4): 972-1059.

Borghans, L., H. Meijers and B. ter Weel. 2008c. "The Role of Noncognitive Skills in Explaining Cognitive Test Scores." *Economic Inquiry*, 46(1): 2-12.

Boschini et al (2011)

Cabrales, A., A. Calvó-Armengol and N. Pavoni "Social Preferences, Skill Segregation, and Wage Dynamics," *Review of Economic Studies*, 2008, 75, 65-98.

- Caselli, F. "Technological Revolutions," *American Economic Review*, 1999, 89(1), 78-102.
- Costinot, A. (2009), "An Elementary Theory of Comparative Advantage," *Econometrica*, 77:4, 1165-1192.
- Costinot, A. and J. Vogel (2010), "Matching and Inequality in the World Economy," *Journal of Political Economy*, 118:4, 747-786.
- Davidson, C., F. Heyman, S. Matusz, F. Sjöholm, and S. Zhu (2011), "Globalization and Imperfect Labor Market Sorting," IFN Working Paper No. 856.
- Davis, S. and Haltiwanger, J. (1991). 'Wage Dispersion within and between US manufacturing plants 1963-1986', in Martin Baily and Clifford Winston, eds., *Brookings papers on Economic Activity, Microeconomics* 1991, 115-200.
- Dunne (1997)
- Dunne, T., Foster, L., Haltiwanger, J., and Troske, K.R. (2004). 'Wage and Productivity Dispersion In US Manufacturing: The Role of Computer Investments', *Journal of Labor Economics*, vol. 22 (2), pp. 397-429.
- Eeckhout, J. and P. Kircher, "Sorting and decentralized price competition," *Econometrica*, 2010, 78(2), 539–574.
- Eriksson, T., M. Pytliková and F. Warzynski (2009), "Increased Sorting and Wage Inequality in the Czech Republic: New Evidence Using Linked Employer-Employee Dataset," mimeo.
- Faggio, G., K. G. Salvanes and J. Van Reenen, "The Evolution of Inequality in Productivity and Wages: Panel Data Evidence," CEP Discussion Paper No 821, 2007.
- Fox, J. (2010), "Estimating Matching Games with Transfers," mimeo, University of Michigan.
- Garicano, L. and E. Rossi-Hansberg (2006), "Organization and Inequality in a Knowledge Economy," *Quarterly Journal of Economics*, 121(4), 1383-1435.
- Grossman, G. M. and G. Maggi (2000), "Diversity and Trade," *American Economic Review*, 90(5), 1255-1275.
- Grossman, G. M. and E. Rossi-Hansberg (2008), "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98(5), 1978-1997.
- Grönqvist, E., B. Öckert and J. Vlachos (2012),
- Harrison, A., J. McLaren and M.S. McMillan (2011), "Recent Perspectives on Trade and Inequality", *Annual Review of Economics*, 3, 261-289.
- Heckman J. J., J. Stixrud, and S. Urzua. 2006. "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior." *Journal of Labor Economics*, 24(3): 411-482.
- Hellerstein, J. K. and D. Neumark (2008), "Workplace Segregation in the United States: Race, Ethnicity, and Skill," *Review of Economics and Statistics*, 2008, 90(3), 459–477.
- Hornstein, A., P. Krusell and G. Violante, "The Effects of Technical Change on Labor Market Inequalities," in *Handbook of Economic Growth, Volume 1B*, edited by Philippe Aghion and Steven N. Durlauf, 2005.
- Håkanson C., E. Lindqvist and J. Vlachos (2012), "Changes in the Distribution of Cognitive Skill: Evidence from the Swedish Enlistment Battery Test 1969-1993", work in progress.

Iranzo, S., F. Schivardi and E. Tosetti, "Skill Dispersion and Firm Productivity: An Analysis with Employer-Employee Matched Data," *Journal of Labor Economics*, 2008, 26(2), 247-285.

Jones, C. (2011), "Intermediate Goods and Weak Links in the Theory of Economic Development", *American Economic Journal Macroeconomics*, 3(2), 1-28.

Katz, L. F. and D. H. Autor (1999), "Changes in the Wage Structure and Earnings Inequality," in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, Amsterdam, North-Holland.

Kremer, M. (1993), "The O-ring Theory of Economic Development," *Quarterly Journal of Economics*, 1993, 551-575.

Kremer and Maskin (1996), "Wage Inequality and Segregation by Skill," NBER Working Paper 5718.

Legros, P. and A. F. Newman (2002), "Monotone Matching in Perfect and Imperfect Worlds," *Review of Economic Studies*, 69, 925-942.

Legros, P. and A. F. Newman (2007), "Beauty Is a Beast, Frog Is a Prince: Assortative Matching with Nontransferabilities," *Econometrica*, 75(4), 1073-1102.

Lentz, Rasmus and Mortensen, Dale T. (2010), "Labor Market Models of Worker and Firm Heterogeneity," *Annual Review of Economics*, 2: 577-602.

Lindqvist, E. and R. Vestman, "The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment," *American Economic Journal: Applied Economics*, 2011, 3(1), 101-128.

Lise, J., C. Meghir and J.-M. Robin (2008), "Matching, Sorting and Wages," mimeo.

Lopes de Melo (2009), "Sorting in the Labor Market: Theory and Measurement," mimeo, University of Chicago.

Machin (1996)

Milgrom, P. and J. Roberts (1990), "The Economics of Modern Manufacturing: Technology, Strategy, and Organization," *American Economic Review*, 80(3), 511-528.

Nordström-Skans, O., P-A Edin and B. Holmlund (2009), "Wage Dispersion Between and Within Plants: Sweden 1985-2000," in *The Structure of Wages: An International Comparison*, E. P. Lazear and K. L. Shaw (eds), University of Chicago Press.

Ohnsorge, F. and D. Trefler (2007), "Sorting It Out: International Trade with Heterogeneous Workers," *Journal of Political Economy*, 115(5), 868-892.

Rosen, S. (1981), "The Economics of Superstars," *American Economic Review*, 71(5), 845-858.

Rosen, S. (1982), "Authority, Control, and the Distribution of Earnings," *Bell Journal of Economics*, 13(2), 311-323.

Saint-Paul, G. (2001), "On the Distribution of Income and Worker Assignment under Intrafirm Spillovers, with an Application to Ideas and Networks," *Journal of Political Economy*, 109(1), 1-37.

Segal, I. and D. Carmit. 2008. "Working When No One is Watching: Motivation, Test Scores, and Economic Success," *Journal of the European Economic Association*, forthcoming.

Shimer, R and L. Smith (2000), "Assortative Matching and Search," *Econometrica*, 68(2), 343-369.

Spitz-Oener, A. (2006), "Technical Change, Job Tasks and Rising Educational Demands: Looking Outside the Wage Structure," *Journal of Labor Economics*, 24(2), 235-270.

9 Appendix

[TO BE INCLUDED]