Network for Studies on Pensions, Aging and Retirement

Netspar THESES

Daniel Pollmann

Risk Preferences and Occupational Sorting

Combining Survey and Top-Coded Administrative Wage Data

MSc Thesis 2011-036

Risk preferences and occupational sorting:

Combining survey and top-coded administrative wage data

Daniel Pollmann*

Master thesis, MSc Econometrics and Operations Research

Supervisors: Thomas Dohmen (ROA, AE2), Franz Palm (KE)

School of Business and Economics, Maastricht University

July 26, 2011

Abstract

We find that the risk aversion of individuals in a household survey (SOEP) is negatively correlated with the wage dispersion in their occupation, as estimated on an administrative wage data set (IABS). Due to top-coding of wages, a three-step censored quantile regression approach is used, while the estimation of occupation-level dispersion relies on the differences of group coefficient estimates at different quantiles.

Keywords: Occupational choice; Sorting; Risk preferences; Earnings risk; Censoring; Quantile regression; Dispersion estimation; SOEP; IABS

*i489212, d.pollmann@student.maastrichtuniversity.nl

Contents

1 Introduction

The primary financial risk faced by many people is labor earnings risk, which varies across occupations. At the same time, individuals differ in their appetite for risk, suggesting that it may affect occupational choice. Indeed, we find that individuals with greater stated willingness to take risks work in occupations with higher cross-sectional wage dispersion.

Through their effect on occupational choice, risk preferences influence agents' utility at the individual level as well as overall efficiency. Particularly in light of a shrinking labor force and shortage of skilled labor in many Western countries, it appears imperative to study the drivers of educational and occupational choice. Workforce composition will be affected by risk attitudes, the value of which, in turn, may vary across occupations.

We rely on two separate data sets. The German Socio-Economic Panel (SOEP) provides us with survey information on risk preferences and other individual and household characteristics. The IAB Employment Sample (IABS) contains administrative wage data collected for social security purposes, which we use to estimate human capital models and compute statistics for wage dispersion within occupations; due to right-censoring of wages, we follow a three-step censored quantile regression approach (Chernozhukov and Hong, 2002). After estimating risk profiles of occupations on the IABS data, we match them to individuals in the SOEP working in these occupations, and analyze the relationship of wage dispersion and individuals' stated willingness to take risks on an eleven-point scale.

In contrast to the main previous work (Bonin et al., 2007), the IABS offers much greater sample size, such that we are able to work with more precise occupation definitions and reduce the impact of aggregation on variation. Furthermore, we use a new method to estimate dispersion, which is based on group coefficient estimates at different quantiles rather than residuals. We thereby circumvent working with censored residuals, which frequently requires parametric assumptions. Our semiparametric approach is generally distribution-free and not only admits, but also allows to estimate, differential patterns of dispersion across occupations. This stands in marked contrast to the popular Tobit estimator, which assumes normality and homoskedasticity, the latter of which is at odds with the principal aim of our analysis.

The organization of the paper is as follows. Section 2 provides a detailed introduction to the literature on the labor market impact of risk preferences, with a particular emphasis on occupational sorting. The two data sets, SOEP and IABS, are described in section 3.

Section 4 discusses estimation of wage regressions on censored data, focusing on semiparametric methods. A new method to estimate dispersion, which we consider particularly suited in the present context, is introduced in section 5. We apply this method to obtain the main results on cross-sectional risk in section 6. Some extensions, such as domain-specific risk and panel estimation, are presented in section 7, while section 8 concludes.

2 Literature on risk and sorting

An early investigation of heterogeneity in risk preferences and occupational sorting can be found in Bellante and Link (1981), who consider the choice between private sector and public sector employment. The authors model this choice in a probit framework, using data from the Panel Study of Income Dynamics (PSID) including a measure of risk aversion based on the quality and insurance of cars owned, the use of seat belts, the extent of medical coverage, and drinking and smoking habits.[1] They find that a high degree of risk aversion significantly increases the probability of public sector employment. While their analysis does not consider the relative riskiness of professions, the authors refer to previous findings which suggest that the probability of unemployment is much lower in the public than in the private sector. Pfeifer (2011) uses SOEP data to find that a greater stated willingness to take risks significantly decreases the probability of public sector employment; the effect is stronger for willingness to take risks in a career context than for general willingness. In his own data based on a student questionnaire, the author additionally elicits preferences for employment security, and finds that this measure has an even stronger effect than both general and career-related risk preferences.[2] Guiso and Paiella (2005) analyze the same question using data from the Bank of Italy Survey of Household Income and Wealth (SHIW); they infer a measure of absolute risk aversion from individuals' willingness to pay for a hypothetical security. Similar to Bellante and Link, they estimate a

[1] This revealed preference approach uses actual choices taken under constraints; hence, these do not only reflect preferences, as noted by Bellante and Link (1981). If the constraints and the dependent variable are correlated, this form of measurement may induce an endogeneity problem. For instance, medical coverage may be cheaper or easier to obtain for public sector employees.

[2] This effect, however, might partially stem from the questionnaire design, as the author alerts participants to the difference in unemployment risk between the public and private sector in his description of the job options available to respondents ("framing"). Furthermore, it is not clear whether the elicitation of preferences for employment security is valid; it is certainly not incentive-compatible. Validation is possible only if one shifts the focus of analysis and views the sorting decision into public sector employment as evidence for validity. For all of the analyses discussed, it is necessary to view risk preferences as given and validly measured; in some cases, they are inferred from real investment decisions or actual/hypothetical lotteries and have immediate economic content. Elsewhere, as in the SOEP, willingness to take risks is directly asked for; the validity of this SOEP measure is discussed in section 3.1.2.

probit model and find that a higher degree of risk aversion significantly increases the probability of public sector employment.

Fuchs-Schündeln and Schündeln (2005) exploit the low degree of risk exposure in the public sector to compare wealth holdings of East German households, for whom they argue the choice of occupation and earnings risk exposure to be independent of risk preferences,[3] and of West German households, for whom the choice of occupation cannot be assumed to be independent of risk preferences. They find that precautionary wealth holdings in the West are much lower than in the East, which they link to self-selection of more risk-averse individuals into less risky occupations in the West.

In addition to their results on private and public sector employment, Guiso and Paiella find that an increase in risk aversion significantly lowers the probability of self-employment. This is in line with the results of Cramer et al. (2002), who find that individuals in the Dutch "Brabant" survey with a high reservation price for a hypothetical lottery are much more likely to be self-employed; for measures of relative and absolute risk aversion constructed from the reservation price, this probability is decreasing in the degree of aversion. As self-employment is typically viewed as more economically risky, this confirms expectations; furthermore, it is in line with the theory of entrepreneurship put forward by Kihlstrom and Laffont (1979), who view self-employment as a risky investment. However, Cramer et al. point out an important problem: Since their elicitation of risk preferences occurred long after the decision on self-employment for the individuals in their sample, possible endogeneity of risk preferences makes any interpretation of causality difficult. Specifically, the experience of self-employment and its connected risks may have had a moderating effect on risk aversion. Ekelund et al. (2005) find that more risk-averse individuals in the Northern Finland 1966 Birth Cohort Study are significantly less likely to be self-employed; since they use a psychology measure, harm avoidance, which develops well before adulthood and is stable from thereon, they give a causal interpretation to their results. On data from the Health and Retirement Study (HRS), Barsky et al. (1997) investigate the predictive power of survey responses to hypothetical future income gambles for selection into self-employment; the relationship is of the expected direction, though very noisy. Beauchamp et al. (2011) analyze data from the Swedish Screening Across the Lifespan Twin (SALT) study,

[3]The authors use the German reunification as a natural experiment. Since labor income risk was largely absent in the GDR, they argue that risk preferences (at least with respect to financial risk) did not play a role in individuals' initial sorting decisions in the GDR. After reunification, certain occupations remained in the public sector with little exposure to income or unemployment risk, while individuals in other occupations were suddenly exposed to income risk.

and find that a risk question in the spirit of that in the HRS is strongly correlated with the probability of having owned a business. Caliendo et al. (2009) as well as Dohmen et al. (2011b) consider stated willingness to take risks in SOEP data and find that individuals with lower degrees of risk aversion are significantly more likely to start a business; however, this only holds for those coming out of dependent employment, and not unemployment or other inactivity. Hartog et al. (2002) use three different Dutch data sets (including the "Brabant" data), and find evidence suggesting that the self-employed are significantly less risk averse, while public sector employees are significantly more risk averse.[4]

Applying an ordered generalized extreme value model on data from the UK Family Expenditure Survey (FES), Brown et al. (2006) analyze the choice between not only fixed-wage contracts and self-employment, but also performance-related pay with an intermediate degree of earnings risk. They find that more risk tolerant individuals, as measured by proxies for risky behavior (expenditures on gambling, smoking and alcohol on the one hand and insurance on the other hand), are more likely to select into the two more risky income settings, the denizens of which are rather similar in their risk-taking characteristics. Using privately collected Italian data including the HRS income gamble question, Di Mauro and Musumeci (forthcoming) obtain the same results regarding the risk preferences of individuals in fixed and variable income employment.

Somewhat startingly, Brunello (2002) uses the same SHIW data as Guiso and Paiella, but comes to very different conclusions:[5] In a probit model, he finds that risk aversion does not affect the probability of public sector employment at conventional significance levels. It is notable that while both papers use the same survey question to infer absolute risk aversion, their interpretation of the lottery implied by this question is different.[6] Furthermore, due to the different nature of their studies, Brunello partials out the effects of financial wealth from his measure of risk aversion, whereas Guiso and Paiella control for financial background in their later regressions. While it is unclear whether these differences are solely responsible for the disparity in results, they nonetheless highlight the difficulty in unambiguously eliciting risk aversion, even just of the financial kind. Brunello additionally fits a multinomial logit model for

[4] Not all results hold true for all data sets, but Hartog et al. provide possible explanations.

[5] Brunello uses risk aversion as an instrument for schooling to estimate the returns to education. For this, he needs to show that risk aversion does not affect income through the channel of occupational sorting. I will discuss several additional points which call into question the validity of Brunello's exogeneity assumption.

[6] Both interpretations, however, imply a degree of absolute risk aversion which is decreasing in the reservation price stated, although the curvature of the transformations is different.

selection into one of ten categories constructed as blue collar workers as well as white collar and managerial employees in each of the industrial, service, and public sectors plus school teachers in the public sector; as before, absolute risk aversion is insignificant at conventional levels.

Of course, individuals do not only choose between dependent employment and entrepreneurship or the private and the public sector; they also choose particular occupations. Saks and Shore (2005) consider a multinomial logit model for the choice between college majors leading into eight different types of occupations.[7] They take major and other student characteristics from the National Postsecondary Student Aid Survey (NPSAS); in addition to some standard controls, this survey also contains information on the wealth of students' parents. The authors show that under decreasing absolute risk aversion utility, individuals with greater wealth should opt for majors associated with riskier occupations. To infer the riskiness of occupations, they estimate the variation left unexplained by human capital models per occupation. They do so in a cross-sectional framework using data from the Baccalaureate & Beyond (B&B) survey, and in a panel setting using PSID data; in the latter specification, the authors can control for unobserved heterogeneity.[8] For each of the eight different fields, Saks and Shore estimate the change in the entrance probability resulting from an increase in wealth and find that those majors which tend to be chosen more frequently at higher levels of wealth are associated with greater occupational earnings risk. This confirms the results of the early analysis of King (1974), who finds that individuals from wealthier families work in occupations with greater income variance, as measured within age-cohort groups. Caner and Okten (2010) carry out a similar analysis for Turkish university entrance exam data (OSS): Students with higher family income or from a family of entrepreneurs are significantly more likely to choose majors perceived as risky; however, the authors cannot estimate the riskiness of occupations, but only reason by their share of public sector employment. Le et al. (2011a) use the Australian Twin Study of Gambling containing a measure of risk attitude similar to the stated willingness to take risks in the SOEP, and estimate a multinomial logit model to find a predictive effect on employment in occupation categories.

The above analyses generally do not consider sorting into occupations, but rather categories of occupation, or, even broader, the choice between dependent employment and entrepreneur-

[7] The types are teachers, healthcare workers, computer workers, engineers, math/science workers, sales people, managers, and entertainers.

[8] Individual-level effects are relevant to this kind of analysis if their variance systematically differs by occupation, and if they are (partially) known to individuals when choosing an education/occupation.

ship. Since, depending on the level of aggregation, common occupational classifications frequently consist of more than a hundred different occupations, it is often not feasible to estimate this choice using a categorical model. Moreover, a mere assignment to occupations does not reveal the motive of risk-based sorting; not only is occupational risk multidimensional (earnings, unemployment, health, etc.), but there may also be unobservable factors, which are rather hard to make out if the relevant risk is not tractable in the first place. However, if it is possible to compute explicit measures for the riskiness of occupations, one can estimate the relationship between an individual's degree of risk aversion and the riskiness of his or her occupation. This is the approach taken in Bonin et al. (2007) and Fouarge et al. (2011) as well as in this paper; a detailed description of the approach will be given in the relevant sections. In all cases, a human capital model for wages is fitted; the former two papers use the dispersion of the residuals within a particular occupation as a measure of riskiness. Bonin et al. infer both risk preferences and riskiness of occupations from responses to questions in the SOEP; individuals are asked to state their general willingness to take risks on a scale and to report their gross monthly income. As expected, individuals who are more willing to take risks work in occupations which are significantly more risky as measured by the dispersion of earnings. Fouarge et al. consider Dutch secondary-school graduates who participated in the ROA School Leavers Survey, which contained the same question on general willingness to take risks as well as information on the graduates' initial jobs. Unlike Bonin et al., they use a separate data set, the Dutch Labor Force Survey, to infer the riskiness of occupations. As risk preferences are elicited around the point of entering the labor force, potential endogeneity problems resulting from exposure to particular occupation characteristics are reduced.[9] It is found that one year after the initial survey round, individuals with high willingness to take risks are indeed employed in more risky occupations,[10] with respect to not only earnings dispersion, but also cyclical sensitivity and unemployment risk. Furthermore, individuals who are initially poorly matched with respect to their preferences and the characteristics of their job are significantly more likely to switch to a more fitting job. Findeisen (2010) exploits the longitudinal dimension of the SOEP to estimate both transitory and permanent income risk within occupations; he finds that more risk tolerant

[9] Individuals' risk preferences are not elicited exactly when they enter the labor market, but slightly later. Furthermore, the occupation-specific risk may affect graduates' risk preferences even before entry; this may occur due to actual exposure in the process of applying for jobs, but even on a "softer" level, while students survey career opportunities and mentally settle into particular roles.

[10] In addition, patient individuals tend to be employed in occupations with steeper wage growth, which may be thought of as delayed compensation.

individuals — measured by an imputed coefficient of constant relative risk aversion (CRRA) — are employed in occupations with greater income risk of both transitory and permanent nature. While in a welfare model, heterogeneous risk preferences indeed yield a better fit than a representative-agent specification, a counterfactual experiment eliminating uninsurable income risk suggests only a minor effect of risk on occupational sorting.

Isphording (2010) replicates the cross-sectional results of Bonin et al. linking occupational choice and risk attitudes in the SOEP to earnings variability computed on German microcensus data. Using a finer level of occupational coding than Bonin et al., he finds effects about three times as large. On the other hand, in a panel specification of occupational choice based on cross-sectional earnings risk and including Mundlak terms, he finds that effects are no longer significant. However, we argue in section 3.1.2 that this interpretation may very well be flawed. In their multinomial model, Le et al. additionally consider a grouping of occupations according to the cross-sectional standard deviation of earnings estimated following the approach of Bonin et al.. They do not arrive at any clear results, and argue that other factors, such as masculinity or femininity of occupations, are more relevant for risky decisions. On the other hand, this grouping of otherwise possibly very dissimilar occupations may distort the relationship of other, perhaps more forceful, factors, such that the effect of risk tolerance cannot be validly estimated either. While not reported in their paper, the earnings measure of their data set, the Australian Twin Study of Gambling, is in categorial form according to income brackets, which will reduce the precision of the variance estimates.[11]

As part of an analysis of risk sharing, Schulhofer-Wohl (2009) uses responses to the HRS risk question to relate them to the amount of income risk experienced by individuals. To estimate the latter, he uses social security earnings records contained in the HRS; due to top-coding of wages, a Tobit model is employed.[12] Schulhofer-Wohl classifies individuals into a low and a high risk tolerance group and finds that the latter carry significantly more of both aggregrate risk (sensitivity to aggregate shocks, cyclicality) and idiosyncratic risk. He points out that since individuals are interviewed late in life, risk preferences may have adjusted to the fluctuations experienced. For a variation of the HRS risk question included in the 1998

[11] The same is true for the microcensus data considered by Isphording, which suffers from the additional complication that total rather than labor income alone is reported.

[12] Schulhofer-Wohl estimates a panel Tobit model, in which he allows heteroskedasticity at the cross-sectional though not longitudinal level. The normality assumption remains crucial, particularly in light of the high degree of censoring and its volatility over time ("30 to 60 percent of observations on prime-age male workers are censored in each year", Schulhofer-Wohl, 2009, p. 10); these points will be intensively discussed at a general level in section 4.2.

wave, Sahm (2008, p. 18) finds that among individuals without prior unemployment spell, those who will experience future unemployment are significantly more risk tolerant, suggesting that "high risk tolerance types have systematically chosen riskier careers with a higher chance of displacement." Guiso et al. (2002) use expectations on income risk and unemployment from the SHIW and find that risk aversion is negatively correlated with the variance of expected income, though not earnings; expected income, here, is calculated given expected earnings from employment, subjective probability of unemployment, and the replacement rate (unemployment benefits).

Several empirical studies document compensation for wage dispersion; Hartog (2009) provides an overview and formally derives a risk-augmented Mincer equation. More subtly and closer to our analysis, groups of individuals might demand different levels of compensation based on their risk aversion. Since women are frequently found to be more risk averse than men (e.g., Barsky et al., 1997; Borghans et al., 2009; Dohmen et al., 2011b),[13] one may expect them to receive higher compensation. Feinberg (1981), however, finds that women receive significantly lower compensation for intertemporal variance than men; in the cross section, the evidence is mixed, from lower or no compensation at all for women (Hartog and Vijverberg, 2007; Berkhout et al., 2010) to equal or even higher compensation (McGoldrick, 1995; Hartog et al., 2003). Similarly, for a given level of risk, public sector workers would be expected to receive higher compensation; however, Berkhout et al. (2010) report evidence of the opposite.

More specifically, there are several ways in which risk preferences may directly correlate with wages. Risk tolerant individuals may sort into (economically) riskier jobs and be rewarded with a risk premium. Leigh (1981) predicts that less risk-averse individuals will work in occupations with higher mean wage and variance; he estimates an employment frontier of occupation-specific mean-variance profiles on the PSID and finds a generally positive relationship of mean and variance. Less risk-averse individuals may also more frequently work in occupations with greater health or mortality risk and receive a compensating wage differential. In addition, it is conceivable that risk aversion, or tolerance, is a productive trait itself.[14,15] Taking a differ-

[13]The support for this argument is often stronger in the field than for experimental data. See Eckel and Grossman (2008) for a survey.

[14]The sign of the effect is not entirely clear and may differ by occupation. While in commercial or managerial careers, a higher propensity to take risks may be rewarded, greater conservatism may be desirable in flight safety officers.

[15]It should be clear to the reader at this point that the estimation of causal effects of risk tolerance on earnings is subject to possibly severe endogeneity problems due to sorting and, as will be discussed below, correlation of risk tolerance and ability; in addition, an interaction of the two factors appears likely. Le et al. (2011b, pp. 558-59, table 3) find an "effect" of risk tolerance on earnings almost twice that of Bonin et al. (2007, p. 930,

ent perspective on wage determination, job-search theory expects more risk-averse individuals to have a lower reservation wage, resulting in a lower mean wage (Pissarides, 1974; Cox and Oaxaca, 1989).[16,17]

If higher education itself is considered risky, this may act as another channel, affecting both first and second moment of the conditional wage distribution. As yet, studies on the relationship of risk aversion and schooling have not led to a consensus, but rather made conflicting predictions; while education may be considered a risky investment (Levhari and Weiss, 1974), it may also shield against unemployment (Mincer, 1991; Nickell and Bell, 1996).[18] Instrumenting schooling attainment by risk aversion, Brunello's first-stage regression yields a significant negative relationship between risk aversion and years of education for Italian SHIW data. Using the slightly different measure of absolute risk aversion of Guiso and Paiella, Belzil and Leonardi (2007) fit a reduced-form dynamic discrete choice model, which predicts that school continuation probabilities decrease with risk aversion at low grade levels, but increase at high levels. However, risk attitudes only explain a modest portion of these probabilities. For the same risk measure, Belzil and Leonardi (2009) incorporate endogeneity of risk aversion and wealth and find a large negative effect of risk aversion on the probability of entering higher education upon completing high school. Belzil and Hansen (2004) assume heterogeneous abilities but homogeneous risk attitudes and estimate a dynamic programming model on data from the National Longitudinal Survey of Youth (NLSY). They find that wage and employment rate dispersion are decreasing with schooling attainments, and that a counterfactual increase in risk aversion will raise (mean) schooling attainments. Chen (2008) presents an empirical framework for estimating wage variance in the presence of selection bias and unobserved heterogeneity, which she implements on NLSY data. She finds a (hump-shaped) effect of education on wage uncertainty, increasing at college entry, but subsiding at completion. While the initial effect may deter risk-averse agents, uncertainty may also reflect an expanded choice set (Weisbrod, 1962). Shaw (1996) estimates a joint model of the investment in financial wealth and human capital on data from the Survey of Consumer Finances (SCF). Using CRRA coefficients inferred from

table 1), but unlike the latter, they do not control for occupation.

[16] With respect to the dispersion of earnings, Pissarides shows that the range of incomes can actually be greater for a more risk-averse society, which contradicts the conjecture of Friedman (1953).

[17] Feinberg (1977) hypothesizes that (1) the expected duration of unemployment is increasing in the variance of potential wage offers, (2) a more risk-averse individual will have a shorter expected duration of unemployment; estimation of a reduced-form equation on PSID data yields supporting results.

[18] See also Guiso et al., who document that subjective unemployment risk, i.e., the expectation of the probability of unemployment, is decreasing in education. Risk attitudes, on the other hand, do no significantly affect subjective unemployment risk.

individuals' investment allocation — thus assuming stable risk preferences across domains — she finds that investment in schooling is decreasing in risk aversion and leads to higher wage growth; as a by-product, more risk tolerant individuals are estimated to experience greater wage fluctuations and variability in wage growth. Shaw (1996, p. 647) suggests that "risk taking is in part an ability to comprehend uncertain outcomes and to make intelligent decisions regarding them."

Related to this, a correlation between risk attitudes and (typically) unobserved individual ability has been suggested; Benjamin et al. (2006), Dohmen et al. (2010), and Beauchamp et al. (2011) present evidence for a negative relationship between risk aversion and cognitive ability for a variety of risk measures and data sets. More risk tolerant individuals may therefore select into occupations offering more meritocratic compensation schemes on the basis of their superior cognitive skills. However, not everyone within an occupation can end up in the upper regions of the distribution, so this really necessitates a first-order location rather than just a pure scale shift to compensate the (potential) "losers" and entice sufficiently many of them to sort into this occupation and stay in it. In particular, (observed) scale will only be valuable if an occupation also contains a portion of individuals with lower cognitive ability. For both risk aversion and productivity observed, Dohmen and Falk (2010, 2011) present experimental evidence that both more risk tolerant and more productive workers are more likely to sort into variable-pay rather than fixed-pay schemes.

Individuals may also receive non-pecuniary compensation for taking on risk, such as in the form of status or prestige.[19] Risk preferences may therefore reflect a desire to be viewed as brave or daring, possibly explaining part of the gender gap mentioned above. More generally, experimental results by Dohmen et al. (2011a), obtained using MRI scanners, suggest that interpersonal comparisons of income are relevant for individuals' utility.

Several points emerge from the existing literature. Risk attitudes are of potential interest in a variety of labor market contexts. At the same time, empirical representations of risk aversion are broader than a single parameter for the curvature of utility functions familiar from microeconomics, and measurement may be noisy. There is similar breadth on the left-hand side of the equation: While discrete choice models are most natural in analyzing sorting into different

[19] Consider the following excerpt from Diamond (2003, p. 108) on the evolution of food production: "Conversely, men hunters tend to guide themselves by consideration of prestige: for example, they might rather go giraffe hunting every day, bag a giraffe once a month, and thereby gain the status of great hunter, than bring home twice a giraffe's weight of food in a month by humbling themselves and reliably gathering nuts every day."

categories of occupation — self-employment or dependent employment, private or public sector, and different occupations — they fail to quantify the risk of interest.

Previous studies (Bonin et al., 2007; Isphording, 2010; Fouarge et al., 2011) relate the cross-sectional standard deviation of earnings within occupations to the risk attitudes of individuals. However, their dispersion measure is superficial in its description of the conditional wage distribution and does not explicitly incorporate the dispersion effect of covariates. With a very large administrative data set, the IABS, we expand on their analysis.

The literature calls for caution regarding the causal interpretation of results. What looks like sorting may be driven by endogeneity resulting from either reverse causality — risk preferences adjusting to risk experienced — or a shared relationship with a third factor — possibly cognitive ability — leading to correlation rather than causation.

3 Data

The analysis will draw on two distinct data sets: a household survey, the German Socio-Economic Panel (SOEP),[20] and an administrative wage data set, the German IAB[21] Employment Sample (IABS). The SOEP contains detailed information on household and individual characteristics, in particular on occupation and risk preferences. We are able to estimate the riskiness of particular occupations on the IABS data, which contains information on wages and human capital variables collected for social insurance purposes.

In the initial replication intended to compare our method to previous approaches, we follow the sample choice of Bonin et al. (2007, pp. 928-29). It is identical to that described below, except that it includes tenured civil servants and individuals from East Germany. Working exclusively with SOEP data, we use it to infer both individuals' risk preferences and the riskiness of occupations.

In our main analysis, we instead use IABS data for the latter estimation; however, we closely align the two independent data samples in terms of their characteristics, using West German[22] men in full-time employment between 25 and 55 years of age. We restrict our sample to West

[20] For an introduction to the SOEP, see Wagner et al. (2007); more information is also available at `http://www.diw.de/en/soep`.

[21] *Institut für Arbeitsmarkt- und Berufsforschung* (Institute for Employment Research of the German Federal Employment Agency), Nuremberg, Germany, `http://www.iab.de/en/ueberblick.aspx`

[22] In the SOEP, we identify individuals as residents of West Germany, whereas the IABS data only provide information on an individual's workplace, such as its being situated in either West or East Germany; we chose to neglect this slight lack of congruence. Furthermore, we do not differentiate between foreigners and German citizens.

Germany for several reasons: First of all, the economic transformation of the East has not yet led to full-fledged convergence in labor market structures and outcomes; we would likely need to carry out the analyses separately, since the availability of jobs and the attractiveness of different occupations, in terms of their risk profile and otherwise, might systematically differ in East and West. More importantly still, those individuals who made their initial sorting decision in the GDR likely did so independently of their risk preferences.[23] To the extent that the choice of occupation is a sticky decision, due to education, training or institutional arrangements, the relationship between risk preferences and riskiness of chosen occupation might not be as strong for these individuals. As a practical econometric issue, top-coding of wages occurs at a different threshold in East Germany due to a difference in pension contribution limits. While our estimation framework allows for such idiosyncracies, incorporation may introduce additional complications. Our focus on men follows the usual rationale: Women more frequently take several years off from work for childbearing and subsequent care, which might limit their choice set and occupational mobility; in addition, they frequently work in types of employment different from the classical full-time job. As a related point, their role in providing family income may differ from that of men, obscuring the sorting relationship. Moreover, women's labor market decisions are likely to be subject to more severe endogeneity than men's.

Since our primary interest lies in sorting within dependent employment, we do not consider the self-employed.[24] As the IABS data do not include tenured civil servants, we omit them from the SOEP sample as well.[24] However, this does not rule out the presence of public sector employees in our samples, as the German public sector distinguishes between tenured civil servants and ordinary employees.[25] While under certain aspects, it may or may not be desirable to exclude all public sector employment, the IABS data, in contrast to the SOEP, do not contain the relevant information. We exclude part-time workers, apprentices, and workers in retirement, as well as those employed in military service or alternative civilian service.

[23] Cf. the institutional features Fuchs-Schündeln and Schündeln (2005) exploit in their analysis.

[24] See section 2 for existing literature on risk attitudes and entrepreneurship as well as public sector employment.

[25] See Federal Ministry of the Interior (2002) for a general introduction to the public service in Germany, and Federal Ministry of the Interior (2009) for an outline of recent reforms on the federal level, along with some statistics.

3.1 Socio-Economic Panel (SOEP)

3.1.1 General information

The German Socio-Economic Panel (SOEP) either provides us with or allows us to infer information on most of the variables typically included in a human capital model, Mincerian or otherwise: years of education, potential experience (calculated as age minus years of schooling minus six), tenure, and occupation. In addition, subjects report their gross monthly labor income. Following Bonin et al. (2007), we drop the bottom and top percentile of the (unconditional) wage distribution.

3.1.2 Risk preferences

Recent waves of the SOEP contain self-reported measures of risk aversion in general and in specific contexts. In particular, individuals were asked to state their willingness to take risks on an eleven-point scale, where zero corresponds to complete unwillingness to take risks, and ten to complete willingness. The general risk question was included in 2004, 2006, 2008, and 2009.[26]

While these self-reported measures are an attractive way of eliciting risk attitudes in a variety of domains in a fast and cost-efficient way, it may not be immediately clear whether the responses reflect true underlying traits. Dohmen et al. (2011b) confirm the validity of the general risk question by analyzing its predictive power for decisions in a Holt and Laury (2002) type lottery with real money, which they administer on a separate sample. In addition, all of the survey measures explain a variety of risky behaviors, including holding stocks, smoking, and participation in active sports. While the domain-specific risk preferences are the best predictors in their specific context, the general risk question fares best as an overall measure. Notably, it also clearly outperforms lottery decisions in explaining risky behavior in non-financial contexts.

It is clear that due to costs arising from, e.g., education, capital accumulation, and search, individuals do not continuously update their choice of occupation, nor likely at a very frequent rate.[27] For current risk aversion and earnings risk to correlate, preferences therefore need to

[26] The exact wording (translated from German) is as follows: "How do you see yourself: are you generally a person who is fully prepared to take risks or do you try to avoid taking risks? Please tick a box on the scale, where the value 0 means: 'not at all willing to take risks' and the value 10 means: 'very willing to take risks'."

[27] Early in their working life, individuals appear more likely to experiment with jobs; later, improved matching and knowledge about productivity (Harris and Weiss, 1984), capital accumulation, a shorter remaining time frame to recoup costs of switching, and background characteristics will likely reduce this behavior.

display a certain degree of stability.[28] This does not rule out deterministic patterns of change, such as the positive age gradient of risk aversion documented in Dohmen et al. (2011b).[29]

Likewise, as pointed out in the literature survey in section 2, any causal sorting interpretation of results rests on some sort of stability assumption with respect to risk attitudes, in particular, exogeneity of the occupational risk experienced over the working life. Ekelund et al. (2005) use a psychology measure, harm avoidance, resulting from a biosocial theory of normal personality variation (Cloninger, 1987), which develops well before adulthood and is stable from thereon;[30] endogeneity is therefore concluded to not be a problem.

Dohmen et al. (2007) analyze the stability of responses to the general risk question in the SOEP. For two subject pools, one a subset of the SOEP, the other separate, they find a test-retest correlation of 0.62 and 0.60, respectively, over a six-week horizon, and attribute all variation to measurement error. The correlation between the 2004 and 2006 waves of the SOEP, in comparison, is 0.50, which is not too far below the six-week benchmark; this suggests that risk attitudes constitute an inherent and stable trait. Furthermore, except for a worsening health status, which leads to an increase in risk aversion, life events do not exert any influence. This is in line with the results of Sahm (2008) for the HRS risk question:[31] While certain individual background characteristics predict gamble responses, most of these are constant per individual over time, and only few, such as aging and macroeconomic conditions, cause systematic change. Beauchamp et al. (2011) and Lönnqvist et al. (2011) show that the SOEP general risk question performs well in terms of test-retest reliability when compared to other risk measures. Beauchamp et al. additionally discuss the role of measurement error and find that it can be substantial.

Isphording (2010) uses data from the 2004 and 2006 waves of the SOEP to extend the analysis of Bonin et al. (2007) to the panel case, similarly regressing the cross-sectional standard deviation of earnings risk within an individual's occupation on risk preferences and covari-

[28] Similarly, labor markets should be characterized by some level of stationarity; this issue is revisited in section 7.3.

[29] Random future events, such as marriage, can be incorporated in individual choice; however, they will reduce the strength of the relationship between risk aversion and earnings risk.

[30] The empirical evidence on the stability of harm avoidance as a personality trait appears to be limited. The results of Cheung (2007) suggest that "HA [harm avoidance, as defined by Cloninger] is a relatively stable personality trait" for patients with late-life depression; however, only group means are considered. For a different personality inventory, also containing a factor scale for harm avoidance, McGue et al. (1993) find a ten-year retest correlation of .635 (30-day retest correlation: .88). An analysis by McCrae (1993) on the well-known *Big Five* (OCEAN) personality traits discusses several studies reporting correlation coefficients of around .70 for different time horizons, and concludes that much of the remaining variation may be due to measurement error.

[31] Barsky et al. (1997) discuss (persistence of) measurement error in the HRS gamble question on risky jobs.

ates. To control for unobserved heterogeneity (in occupational choice), he employs a Mundlak transformation (Mundlak, 1978), which assumes that individual-level effects are a function of individual time averages of time-variant independent variables. Since the risk preference variable of the SOEP is not time-constant (preferences were separately elicited in 2004, 2006, and later), its individual time average is included as a Mundlak term. In this specification, the coefficient of the regular risk preference variable becomes insignificant, while that of the Mundlak term is positive and large. Isphording concludes that all effects can be attributed to unobserved heterogeneity, in particular ability, with which risk preferences have been found to be correlated. This interpretation is subject to our earlier caveat that (observed) scale is only valuable to high-ability individuals (with possibly low degrees of risk aversion) if sufficiently many individuals of low or average ability (with corresponding greater degrees of risk aversion) have sorted into the same occupation.

More importantly, such an interpretation is misguided if true risk preferences are actually stable, even just over the two-year horizon. In this case, the responses to the SOEP risk question may be viewed as measurements of true latent preferences, with all (non-deterministic) individual-level variation due to measurement error. If the latter averages to zero over time, the individual-level time average is simply a more accurate statistic, which explains why the Mundlak term, rather than the singular observation, captures most of the variation. In the latent framework, this remains true even for a moderate degree of instability, as long as the variance of the measurement error exceeds the within-subject variation in true preferences. Furthermore, due to the stickiness of occupation decisions, the time average of preferences might be the more relevant factor for sorting, at least if the distribution of preferences is somewhat stationary.

Previous findings on independent data sets inspire confidence in the assumption of stable risk preferences. However, measurement error appears to be substantial; while it may exhibit a certain degree of temporal persistence, we nonetheless expect to improve on the measurement of risk preferences by utilizing the responses from multiple waves. For each individual, we therefore compute an unweighted average of all available responses, the number of which varies between one and four.

3.2 IAB Employment Sample (IABS)

Information on the riskiness of professions is inferred from the German IAB Employment Sample (IABS), an anonymized sample from a rich administrative data set. It includes (gross) wage

information as well as other employee characteristics reported by the employer for social security contribution purposes. In this paper, we use the IABS R04 version,[32] which is a 2% sample of the German social security records for the period from January 1, 1975, to December 31, 2004.[33,34]

The IABS is particularly rich because of both its longitudinal (30 years) and its cross-sectional dimension (a total of 1,360,949 individuals). Furthermore, the wage information is highly precise: Since misreporting of wages is subject to severe penalties, measurement error is likely to be minimal; in addition, the IABS comes in the form of spell data such that daily information is available for the entire sample period. From this, we draw a cross section for June 30[35] for each year considered in our analyses.

As described above, we consider West German men in full-time employment between the ages of 25 and 55. We additionally exclude observations with an income reported below the marginal part-time income threshold, which was EUR 400 in 2004, the main year of our analysis. First of all, it appears exceedingly unlikely that these individuals were truly in full-time dependent employment, and not part-time workers, apprentices, in military or alternative civilian service, or in any other status not relevant to our analysis. Second of all, before 1999, incomes below this threshold were not recorded; this exclusion therefore serves the consistency of intertemporal comparisons.

The reporting precision of some of the dependent variables in our human capital model may be considerably lower, since they are only collected and reported for statistical purposes, but with no pertinence to social security. The education variable in particular is frequently missing or inconsistent for different employment spells of one individual. Since measurement error may affect our analysis by introducing possibly systematic noise to our dispersion estimates, we apply a correction described in Fitzenberger et al. (2006). Specifically, we use imputation rule 1, which assumes that there is no over-reporting of degrees and, broadly speaking, extrapolates the highest attained degree forward, and, in case of missing data, also backward.[36] Tenure with

[32] Basic information is available at `http://fdz.iab.de/en/FDZ_Individual_Data/IAB_Employment_Samples/IAB_Employment_Samples_Outline/IAB_Employment_Samples_Regional_File_75-04.aspx`; see Drews (2008) for more detail (in German).

[33] The sample observes proportional stratification of employees in West and East Germany as well as German and foreign citizens.

[34] The data are available from the Institute for Employment Research as a Scientific Use File in two different versions: the regional file used here, and in a weakly anonymized version with more disaggregate geographical and occupational information.

[35] As the midpoint of the year, June 30 is commonly used when drawing cross sections from IABS data (e.g., Ludsteck, 2008; Dustmann et al., 2009; Lehmer and Möller, 2010).

[36] This procedure may result in education levels which are too high. Fitzenberger et al. (2006) propose two

an employer, on the other hand, can be computed with great precision due to the spell nature of the data.

There are two major caveats, the one of a more economic, the other of a more econometric nature, with regard to the data. In Germany, tenured civil servants (*Beamte*) are not insured through the regular social security system, and therefore, as mentioned above, they are not included in our data set. Since the German public sector distinguishes between tenured civil servants and ordinary employees, our data set may nonetheless contain public sector workers of the latter category, who, however, are not identifiable.

From an econometric perspective, a greater challenge arises from top-coding of wages. For the entire sample period, a statutory limit was in place on the amount of monthly income subject to social security contributions (*Beitragsbemessungsgrenze*). This ceiling is a matter of federal legislation and is adjusted on an annual basis; for the main year of our analysis, 2004, it was EUR 61,800, EUR 5,150, and EUR 168.85 for annual, monthly, and daily income, respectively.[37] While it also varies by the location of the employer in either West or East Germany, this is irrelevant for our analysis since we restrict our sample to the former.[38] The degree of censoring remains fairly constant over the sample period, but varies considerably by groups of age, education, and occupation (see table 1).

3.3 Matching of occupations in SOEP/IABS

Both SOEP and IABS report occupations using the German *Klassifizierung der Berufe* ("Classification of occupations"), albeit in different versions: While the SOEP employs the 1992 version (KldB 92), the coding in the IABS is based on the earlier 1988 version (KldB 88).[39] Specif-

other, more conservative rules, but find that both the (marginal) distribution of education and wage regression estimates are not very sensitive to the type of rule used.

[37]While the contribution limit is applied to monthly income, the earnings variable in the IABS data set reports daily income over the employment spell considered, and is censored at the daily contribution limit. For a detailed overview table of contribution limits and part-time income thresholds, see `http://doku.iab.de/fdz/Bemessungsgrenzen_de_en.xls`.

[38]In addition, a different ceiling applies to miners, who are separately insured through the *Knappschaftliche Rentenversicherung* (German Miners Pension Insurance), whereas workers and employees are combined in the *Rentenversicherung der Arbeiter und Angestellten* (Pension insurance of workers and employees). However, the IABS data artificially recensor all miners at the common threshold, which is fine, since the contribution limit for miners has always been above it.

[39]The KldB 88 was developed by the *Bundesagentur für Arbeit* (Federal Employment Agency). More information (in German) is available at `http://statistik.arbeitsagentur.de/Navigation/Statistik/Grundlagen/Klassifikation-der-Berufe/KldB1975-1992/KldB1975-1992-Nav.html`. The KldB 92, based on the KldB 88, was developed by *Statistisches Bundesamt Deutschland* (Federal Statistical Office), with more information (in German) available at `http://www.destatis.de/jetspeed/portal/cms/Sites/destatis/Internet/DE/Content/Klassifikationen/GueterWirtschaftklassifikationen/KlassifikationKldb92,templateId=renderPrint.psml`. Both codings are set to be replaced by a new version, KldB 2010. In contrast to the commonly used international ISCO classification, which is based on skill levels and tasks, the KldB adheres to

ically, the IABS uses the 3-digit level of the KldB 88 (*Berufsbezeichnungen*) identifying 328 different occupations; in our IABS R04 file, these are aggregated to 130 different occupations for anonymity purposes (Drews, 2008, pp. 79-86). In the SOEP, occupations are coded according to the 4-digit level of the KldB 92. To match occupations in the SOEP to those in the IABS, we use a cross walk file from the Federal Statistical Office.[40] Unfortunately, the mapping, which uses the 4-digit level of both versions of the KldB, is not isomorphic; in fact, the correspondence is not unique in either direction. While some of the ambiguous cases are resolved in the inevitable aggregation necessary to conform with the cruder coding of the IABS R04, in many cases, there was no alternative to *ad hoc* assignments. Overall, about two thirds of the individuals in the SOEP sample used for estimation had a non-unique occupation correspondence in the IABS data. While this may sound dramatic, it merely reflects the fact that complex roles and tasks cannot always be easily grouped into occupations; even at the individual level, there will inevitably be a certain degree of noise involved in individuals' reported occupation. As long as the occupation matched in the IABS is representative of the risk profile individuals associate with a particular occupation in the SOEP, this will not be an issue (manual resolution of the ambiguous cases typically yielded rather clear assignments, with much of the prior ambiguity likely due to very special and idiosyncratic correspondences in the cross walk file). It would be a much more severe problem if very dissimilar occupations were grouped together; the coding used, however, distinguishes 130 different occupations and should hence offer sufficiently fine discrimination (in fact, as described above, the finer we choose the coding in the IABS, the more non-unique matchings there will be). The replication study in section 6.1 briefly returns to the coding issue.

4 Censored wage regressions

Consistent estimation of our wage determination model, described below in section 4.1, requires methods which appropriately deal with the top-coding of wages in the IABS data. Section 4.2 discusses classical parametric methods and their limitations, while section 4.3 argues for the use of semiparametric methods and describes estimation.

narrow occupation definitions.

[40] To our knowledge, this file is not publicly available, but it is sent out if requested via their website: `http://www.destatis.de/jetspeed/portal/cms/Sites/destatis/Internet/EN/Navigation/TopNav/Kontakte.psml`

4.1 Mincerian human capital model

To model wages, we choose a log-linear specification in the vein of Mincer (1958, 1974). While these two models are similar algebraically, their economic content is quite different, with consequences for our analysis. Mincer (1958) assumes a homogeneous population induced to undertake schooling investments by compensating differences due to different education requirements across occupations; the log-linear form with respect to earnings and schooling arises from an approximation to earnings maximization. The accounting identity model of Mincer (1974) is less clearly tied to maximization behavior, though it approximates the dynamic human capital investment model of Ben-Porath (1967); the inclusion of covariates may reflect assumptions on the investment ratio and other factors. Here, individuals are *ex ante* assumed to be heterogeneous, with differing returns to schooling; the schooling coefficient in a linear model should therefore only be viewed as an *ex post* average. Perhaps more fittingly, the accounting identity model suggests a random coefficient model. We extend the heterogeneity to the returns to particular occupations and approximate the model at different quantiles using quantile regression.

We include a linear term for years of education by converting attained degrees into years.[41] An alternative to the classical log-linear specification is the inclusion of dummy variables for different educational degrees, which may be a more meaningful measure of educational attainment for Germany (Card, 1999, p. 1806)[42] and is therefore used to gauge the impact of possible misspecification (see section 6.2.3).

The accounting identity model has implications for the experience pattern of variance. Specifically, it suggests a U-shaped pattern around an overtaking age at which earnings variance is due exclusively to differences in schooling levels and ability, but unrelated to post-school investment behavior; before and after, the latter will be an additional source of dispersion.[43] This not only makes estimation at different quantiles appear particularly attractive, but it also suggests a flexible experience specification. We construct (potential) experience as age minus years of education minus six and include a cubic polynomial. A quadratic polynomial is used to incorporate effects of tenure.

The inclusion of occupation dummies serves two purposes: First of all, it captures occupation-

[41] See section C in the appendix for details.

[42] German students select into different streams of secondary school, providing access to different forms of vocational training and higher education such as technical college and university. See KMK (2010) for detailed information on the education system.

[43] For (generally supportive) empirical evidence, see Heckman et al. (2003, 2006) for US data, and Polachek (2003) for US and international data, unfortunately not including Germany.

specific effects such as compensating wage differentials. More central to our analysis, the estimation of occupation effects at different quantiles is useful to evaluate wage dispersion at the occupation level; this point is extensively treated in section 5. In our SOEP replication, we follow Bonin et al. (2007) and include individual risk attitudes as well as dummies for public sector employment and residence in East Germany.

Since we do not intend to estimate returns to schooling or find any associated causal link, we neglect endogeneity bias arising from self-selection into different education streams and unobserved heterogeneity. Rather, we are interested in the earnings variation in an occupation observable to an individual.[44]

4.2 Parametric methods

The Tobit estimator introduced by Tobin (1958) has seen such pervasive use in censored regression that models involving censoring are often generally referred to as Tobit models, even if only loosely associated with the original estimation framework.[45] Tobin (1958) applies the estimator to what should, in economic terms, be more adequately referred to as a corner solution problem, but which is equivalent to the type of censoring problem studied here. The typical model assumes normality and homoskedasticity and can be solved either by using the Tobin-Amemiya maximum likelihood estimator, consistency and asymptotic normality of which are proved in Amemiya (1973), or the two-step Heckit approach introduced by Heckman (1976, 1979).

There is a well-known theoretical literature on deviations from normality (e.g., Arabmazar and Schmidt, 1982; Goldberger, 1983) and homoskedasticity (e.g., Maddala and Nelson, 1975; Hurd, 1979; Arabmazar and Schmidt, 1981; Brown and Moffitt, 1983) in a variety of situations. The biases these authors derive depend on the exact set-up and the degree of censoring, but can be substantial. While the overall degree of censoring in our data is not that extreme compared to their settings, it is within certain occupation cells. The simulation study of Vijverberg (1987) for the case of non-normality shows that the bias can be quite large; of particular relevance to our dispersion analysis is his finding that the estimated error variance is often seriously biased, especially in the case of skewed distributions. Powell (1986b) carries

[44] The observed scale is not necessarily equal to the true scale, which may be truncated from below: Initial hurdles, such as education and certification requirements, may prevent entry; some individuals may not find employment within their desired occupation and decide to take up other employment; others may remain unemployed. In this analysis, we abstract from unemployment risk; rather, we focus on observed earnings variation, which should form part of the agents' information set, much more clearly than any counterfactual variation (see, e.g., Chen, 2008).

[45] For an excellent exposition and discussion, see Amemiya (1984) as well as Amemiya (1985, ch. 10).

out simulations for heteroskedasticity as well as non-normality and finds that "failure of the homoskedasticity assumption may have more serious consequences than failure of normality in censored regression models." Also for the two-step Heckit estimator, simulation evidence suggests that non-normality and heteroskedasticity lead to bias (Paarsch, 1984; Donald, 1995).

The assumption of homoskedasticity contravenes the principal aim of our analysis: to investigate whether more risk tolerant individuals sort into occupations with greater earnings dispersion. Not only would it not be very elegant to operate under such contradictory premises in the two stages of our analysis, but it might also skew our results towards non-significance. Log-normality of wages is frequently assumed, and simple histograms suggest that it fits many data sets, including ours, fairly well. On the other hand, skewness in wages has been documented and discussed (e.g., Lazear and Shaw, 2009), and skewness affection has been found to be a factor in wage determination (Hartog et al., 2003; Hartog and Vijverberg, 2007). Deviations from normality extend to even higher orders: In an analysis on wage discrimination, Chay and Honoré (1998) argue with respect to the fourth moment that "abnormally long tails in the log-earnings distribution is the major source of misspecification in the maximum likelihood estimates of the black-white earnings gap."[46] Since the ML estimator has an unbounded influence function, it is particularly sensitive to such outliers. As any parametric assumption, normality imposes restrictions on shape which we may want to avoid to allow for more differential patterns across occupations and a more intricate choice behavior by individuals.

There have long been tests for the assumptions of normality and homoskedasticity (Chesher and Irish, 1987; Pagan and Vella, 1989), typically taking the form of Lagrange multiplier (LM) tests and using generalized residuals[47] from a Tobit regression. Here, we use two outer-product-of-the-gradient (OPG) versions of the conditional moment tests (CMT) given in Skeels and Vella (1999).[48] The corresponding statistics can easily be computed from an artificial regression (Newey, 1985, p. 1052) and follow $\chi^2(r)$ distributions under the respective null hypothesis of normality or homoskedasticity. Specifically, for the test of normality, $r = 2$; for the test of homoskedasticity against the alternative of multiplicative heteroskedasticity, $\sigma_i^2 = \sigma^2 \exp(\tilde{X}_i'\alpha)$, $r = K + J + 1$, where K is equal to the number of regressors in $X_i \supset \tilde{X}_i$, and $J = \dim(\tilde{X}_i)$. It is is well known that OPG tests can suffer from potentially severe size distortions in small

[46] Cf. the results of Wilhelm (2008, p. 574), who investigates the performance of Tobit, CLAD, and symmetrically censored least squares (SCLS) under deviations from homoskedasticity and normality and finds that "Cauchy errors — symmetric but with extremely long tails — [...] cause the most serious Tobit bias."

[47] On generalized residuals, see, e.g., Gourieroux et al. (1987) and Cameron and Trivedi (2005, pp. 289-90).

[48] See Cameron and Trivedi (2009, pp. 535-38) for implementation in Stata.

samples (Cameron and Trivedi, 2005, pp. 240-41); Skeels and Vella (1999) find partial evidence for this in a Monte Carlo study on the tests described above. Although our sample is large,[49] we rely on a parametric form of the bootstrap (Drukker, 2002) to gain additional confidence about the significance of our test statistics.

A limitation of the two conditional moment tests for the assumptions of normality and homoskedasticity is that they require the respective other to hold; if both tests yield a rejection, we can only conclude that one of the assumptions fails at the desired level. This shortcoming can be addressed using a nested test as in Melenberg and van Soest (1996). The authors further discuss ML estimation of a Tobit model with multiplicative heteroskedasticity or non-normality, the latter using a generalization proposed by Ruud (1984).

4.3 Semiparametric methods

The preceding section has pointed out the limitations of parametric methods with respect to both their consistency under deviations from assumptions and the flexibility of our empirical analysis. Section 4.3.1 introduces (general) semiparametric methods based on quantile regression, which operate under much weaker assumptions. Section 4.3.2 extends these to the case of censoring and discusses estimation methods, in particular the (relatively) recent three-step approach of Chernozhukov and Hong (2002).

4.3.1 Quantile regression

In linear statistical models, one typically relies on least squares methods to compute estimates of conditional expectations. As such, the resulting regression curve provides a summary of the conditional averages. However, if the interest goes beyond levels and their means, the picture will necessarily be rather incomplete. In our analysis, for example, the primary focus is on the dispersion of wages, and we are therefore interested in an adequate statistic. Quantile regression (QR), introduced by Koenker and Bassett (1978) as a generalization of median regression, allows us to parsimoniously describe the entire conditional wage distribution by estimating conditional quantile functions (CQF).[50] In the words of Angrist et al. (2006, p. 539), "unlike

[49] Put less modestly, there will hardly ever be a microeconometric empirical setting in which one may feel closer to the elusive concept of infinity.

[50] An excellent non-technical introduction with illustrative examples and an overview of applications can be found in Koenker and Hallock (2001). Buchinsky (1998) and Angrist and Pischke (2009, ch. 7) summarize a range of points relevant to the empirical researcher, including several which are of interest in our analysis. For a very comprehensive treatment of QR, consult Koenker (2005).

OLS coefficients, QR estimates capture changes in distribution shape and spread, as well as changes in location."

While in linear models, by least squares, we obtain an estimate of the conditional expectation function $\mathbb{E}(Y_i|X_i)$ for $i \in \{1, \ldots, N\}$, QR provides us with an estimate of the CQF $Q_\tau(Y_i|X_i)$, to be defined as follows.

Consider the conditional distribution of Y_i given $X_i \in \mathbb{R}^p$:

$$F_{Y_i|X_i}(r) \equiv Pr(Y_i \leq r|X_i) \tag{1}$$

Then, we denote the conditional τ-quantile of Y given X as

$$Q_\tau(Y_i|X_i) \equiv q_\tau(X_i) \equiv F^{-1}_{Y_i|X_i}(\tau) = \inf_{r \in \mathbb{R}} \{r : F_{Y_i|X_i}(r) > \tau\} \tag{2}$$

It is instructive to consider the corresponding Skorokhod representation:

$$Y_i = F^{-1}_{Y_i|X_i}(U_i),\ U_i|X_i \sim \mathcal{U}(0,1) \tag{3}$$

A semiparametric alternative to linear regression can be based on the Skorokhod representation of a linear quantile model, $q_\tau(X_i) = X_i'\beta(\tau)$, which suggests a random coefficients interpretation:

$$Y_i = X_i'\beta(U_i),\ U_i|X_i \sim \mathcal{U}(0,1) \tag{4}$$

We may assume the linear quantile model to represent the true DGP; however, as will be discussed later, QR also has useful approximation properties. Rewriting the above,

$$Y_i = X_i'\beta(\tau) + u_i(\tau) \tag{5}$$

We do not impose any assumptions on the distribution of the error term; while specifying a parametric model of the conditional quantiles, we are agnostic regarding the error distribution in our semiparametric framework. All we rely on is a conditional quantile restriction, stipulating that the conditional quantile of the error is equal to a constant. We assume that X_i always

includes a constant term for the intercept, which affords us the following normalization:

$$Q_\tau(u_i(\tau)|X_i) = 0 \tag{6}$$

The interpretation of the error term $u_i(\tau)$ is different from the classical model; its stochastic properties derive from the model definition (possibly as an approximation) rather than any prior behavioral assumption.[51] This justifies our interest in the estimation of different quantiles using the methods described in the following.[52]

Together with the assumption that the error has no point mass at its conditional quantile, the conditional quantile restriction leads to a conditional moment restriction:[53]

$$\mathbb{E}[\mathbf{1}(Y_i \leq X_i'\beta(\tau)) - \tau|X_i] = 0 \tag{8}$$

However, estimation is generally simpler when characterizing conditional quantiles as the solution to a particular expected loss minimization problem, in the context of which it is useful to define the "check"[54] (or weighted absolute loss) function. For $\tau \in (0,1)$,

$$\rho_\tau(u) \equiv [\tau \mathbf{1}(u \geq 0)u + (1-\tau)\mathbf{1}(u < 0)(-u)] \tag{9}$$

$$= [\tau - \mathbf{1}(u < 0)]u \tag{10}$$

Quantiles minimize expected weighted absolute loss in a conditional sense:

$$q_\tau(X_i) = \arg\min_q \tau\mathbb{E}[(Y_i - q)^+|X_i] + (1-\tau)\mathbb{E}[(q - Y_i)^+|X_i] \tag{11}$$

$$= \arg\min_q \mathbb{E}[\rho_\tau(Y_i - q)|X_i] \tag{12}$$

In the case of our linear quantile model, $q_\tau(X_i) = X_i'\beta(\tau)$,

$$\beta(\tau) = \arg\min_b \mathbb{E}[\rho_\tau(Y_i - X_i'b))|X_i] \tag{13}$$

[51] Cf. the discussion of nonparametric estimation in Powell (1994, p. 2446).

[52] See Powell (1994, pp. 2469-74) for a discussion of identification via conditional quantile restrictions and associated asymptotic properties.

[53] As a consequence, $\beta(\tau)$ also solves moment conditions of the general form

$$\mathbb{E}\{h(X_i)[\mathbf{1}(Y_i \leq X_i'\beta(\tau)) - \tau]\} = 0 \tag{7}$$

[54] When drawn, the weighted absolute loss function is reminiscent of a check mark.

We can define the QR estimator as its sample equivalent and the optimal predictor minimizing the realized loss:[55]

$$\hat{\beta}(\tau) = \arg\min_b \sum_{i=1}^{N} [\rho_\tau(Y_i - X_i'b)] \tag{16}$$

Typically, (16) is reformulated as a linear programming (LP) problem and solved by simplex or interior point methods. For $\tau = .5$, QR results in estimation of the conditional median and minimization of the sum of absolute deviations. This approach is known as the least absolute deviation (LAD) method:

$$\hat{\beta}_{LAD} = \arg\min_b \sum_{i=1}^{N} |Y_i - X_i'b| \tag{17}$$

Asymptotic normality and consistency of the QR estimator can be shown (Koenker and Bassett, 1978; Koenker, 2005). It should be noted, however, that the estimators of tail quantiles may follow non-standard distributions (Chernozhukov, 2005).

In a Gaussian (or other parametric) *iid* framework, MLE is always most efficient, and QR comes with a loss in efficiency.[56] However, unlike MLE, QR is robust to deviations in distribution (distribution freeness) and heteroskedasticity. Furthermore, since unlike ML estimators, QR estimators have bounded influence functions, they are also more robust to outliers.

Another particular feature, which distinguishes conditional quantiles from conditional expectations, is equivariance to monotone transformations (Koenker, 2005). For any non-decreasing function $h(\cdot)$,

$$Q_\tau[h(Y_i)|X_i] = h[Q_\tau(Y_i|X_i)] \tag{18}$$

[55]Note that $\hat{\beta}(\tau)$ is an M-estimator. This becomes clear if we define

$$Q_n(\beta) = -\frac{1}{n}\sum_{i=1}^{N} [\rho_\tau(Y_i - X_i'b)] \tag{14}$$

which allows us to write

$$\hat{\beta}_n(\tau) = \arg\max_b Q_n(b) \tag{15}$$

[56]The Laplace distribution is an exception. In this case, the LAD estimator (17) is also the ML estimator.

4.3.2 Censored quantile regression

Owing to the above equivariance property, QR is particularly suited for censoring problems.[57] Recall that in the case of top-coding, we observe $Y_i = \min(C_i, Y_i^*)$, where Y_i^* is the latent true value of the process of interest and C_i is some upper limit, of which we assume Y_i^* to be conditionally independent. Since for any $C_i \in \mathbb{R}$, $\min(C_i, \cdot)$ is a non-decreasing function, we have (Powell, 1986a):

$$Q_\tau(Y_i^*|X_i) = X_i'\beta(\tau) \Rightarrow Q_\tau(Y_i|X_i) = \min(C_i, X_i'\beta(\tau)) \tag{19}$$

The censored quantile regression (CQR) estimator follows trivially as the minimizing argument of the Powell objective function (Powell, 1986a):

$$\hat{\beta}_{CQR}(\tau) = \arg\min_b \sum_{i=1}^{N} \{\rho_\tau[Y_i - \min(C_i, X_i'b)]\} \tag{20}$$

Consistency and asymptotic normality can once again be shown (Powell, 1986a). Analogous to before, we have the special case of the censored median regression leading to the censored least absolute deviation (CLAD) estimator (Powell, 1984):[58]

$$\hat{\beta}_{CLAD} = \arg\min_b \sum_{i=1}^{N} |Y_i - \min(C_i, X_i'b)| \tag{21}$$

As in the uncensored case, the QR-based estimator is consistent under general non-normal distributions and heteroskedasticity (Powell, 1984, 1986a). Monte Carlo evidence by Paarsch (1984) suggests that the CLAD estimator (21) has particularly good properties in large samples, with strong performance especially compared to the two-step Heckit, but also the ML estimator.

Although (20) is not an LP problem,[59] the CQR estimate can be obtained by LP methods as in the uncensored case, albeit at greater computational cost. Buchinsky (1994) gives a well-known application to changes in the US wage structure. For estimation, he presents his iterative linear programming algorithm (ILPA), involving an iteration of the Barrodale-Roberts algorithm

[57] An excellent non-technical introduction with an application to wage data can be found in Chay and Powell (2001).

[58] Powell (1984) shows that solving (21) is equivalent to LAD estimation (17) on those observations for which $X_i'\hat{\beta}_{CLAD} < C_i$.

[59] Paarsch converts the problem into a nonlinear program, which he solves by grid search over the parameter space to avoid problems associated with numerical optimization; these problems may arise from the existence of several local optima due to nonconvexity. He comments on both the computational difficulty and the excessive computation time compared to ML and Heckit estimator.

originally developed for LAD regression; ILPA iteratively performs QR on observations with predictions in the uncensored region, based on the previous iteration. Convergence is achieved if two subsequent sets of observations are the same; while this need not occur, convergence guarantees local optimality.[60,61] Another alternative, the BRCENS algorithm, is proposed in Fitzenberger (1997a). Unfortunately, both algorithms have less than reliable convergence properties with respect to the Powell estimator (20), particularly in large samples and for high dimensionality (Chernozhukov and Hong, 2002), as in our application.[62]

Due to the computational difficulty associated with many CQR estimators as well as their somewhat unreliable convergence properties, CQR so far has enjoyed nowhere near the ubiquity in terms of applications one would expect from its flexibility and the prevalence of censoring in empirical analyses. In response, Chernozhukov and Hong (2002) present the three-step CQR method, which avoids a great deal of problems by selecting a more "benign" sample based on an initial regression of the probability of censoring, and subsequently works with standard QR.[63,64]

Step 1. Specifically, let $\eta_i = \mathbf{1}(Y_i \neq C_i)$; that is, η_i is an indicator of non-censoring (with censoring point C_i). We consider a parametric (e.g., probit or logit) model for the probability of non-censoring:

$$\eta_i = p(X_i'\gamma) + \epsilon_i \tag{22}$$

In general, model (22) will be misspecified and therefore inconsistent for the true propensity score $h(X_i, C_i)$. However, it is only used as an auxiliary regression to select an initial sample J_0 with propensity score $h(X_i, C_i) > \tau$, necessary for consistent estimation of quantile τ.[65] To ensure this, we do not base our selection on the condition that $p(X_i'\hat{\gamma}) > \tau$, but rather that $p(X_i'\hat{\gamma}) > \tau + k$, where k is a trimming constant strictly between 0 and $1 - \tau$.

[60] This requires the assumption that the exact CQR estimate does not interpolate any censored observation (Fitzenberger, 1994). For a brief explanation, the reader is referred to Fitzenberger (1997a, pp. 176-77) and Fitzenberger (1997b, pp. 423-24); for the interpolation property, see also Koenker (2005, pp. 33-34).

[61] Buchinsky (1994, p. 412) suggests trying out different starting values for the coefficient vector to initialize the algorithm. If all lead to the same result, "this is an indication that a global minimum is obtained."

[62] Indeed, attempts at regressions using the ILPA algorithm as implemented in the Stata command *clad* (Jolliffe et al., 2000) proved practically infeasible.

[63] Applications include Ludsteck and Haupt (2007) on the responsiveness of wages to changes in labor demand, Gustavsen et al. (2008) on ice cream purchases, Machado and Santos Silva (2008) on pension plan participation, Kowalski (2009) on medical expenditure, Fack and Landais (2010) on charitable giving, and Schmillen and Möller (2010) on lifetime unemployment.

[64] The estimators of Buchinsky and Hahn (1998) and Khan and Powell (2001) similarly carry out a first-stage selection, but are impractical for high dimensionality and large data sets.

[65] Note the deviation from Chernozhukov and Hong, who select a sample with propensity score $h(X_i, C_i) > 1 - \tau$. This is an important difference between left- and right-censoring.

Since we do not necessarily have to select the largest subset J_0, there is some freedom in choosing k. For this, we write J_0 as a function of k, $J_0(k) = \{i : p(X_i'\gamma) > \tau + k\}$. The approach taken here, mirroring that of Chernozhukov and Hong, is to choose the trimming constant k such that

$$\#J_0(k)/\#J_0(0) = 90\% \tag{23}$$

This means that we discard 10% of those observations with a probability estimate higher than τ. Chernozhukov and Hong provide the sufficient condition that $p(X_i'\gamma_0) - k$ (for $\gamma_0 = \operatorname{plim}\hat{\gamma}$) be a lower bound on $h(X_i, C_i)$. For the proof of asymptotic properties, this envelope restriction can be replaced by a much weaker separating hyperplane restriction.

Step 2. In the second step, we obtain the initial (inefficient) estimator of $\hat{\beta}_0(\tau)$ by standard QR on the sample J_0:

$$\hat{\beta}_0(\tau) = \arg\min_b \sum_{i \in J_0(k)} [\rho_\tau(Y_i - X_i'b)] \tag{24}$$

Next, we select $J_1 = \{i : X_i'\hat{\beta}_0(\tau) < C_i - \delta_n\}$, where δ_n is a small number such that $\sqrt{n} \to \infty$ and $\delta_n \searrow 0$. We choose δ_n similarly to k, but with a lower percentage of discarded observations of 3%. The aim of this step is to include all observations $\{X_i : X_i'\beta(\tau) < C_i\}$ to build up the efficiency of the next step.

Step 3. We obtain the three-step estimator $\hat{\beta}_1(\tau)$ by running (24) with J_1 instead of J_0.

Step 4. (optional). Step 3 may be repeated a finite number of times on a sample $J_l = \{i : X_i'\hat{\beta}_{l-1}(\tau) < C_i - \delta_n\}$, yielding estimates $\hat{\beta}_l(\tau)$ for $l \in \{2, \dots\}$.

Chernozhukov and Hong prove consistency of their three-step estimator $\hat{\beta}_l(\tau)$. Due to distributional equivalence with the Powell estimator, the estimator inherits its asymptotic normality and efficiency properties.[66] The conventional variance-covariance matrices for *iid* errors do not extend to the case of independent heteroskedastic observations: What Chernozhukov and Hong (2002, p. 876) refer to as "notational cost" implies a subtle modification with, however, nontrivial consequences for estimation and inference.[67]

[66] In fact, since three-step CQR relies on standard QR methods, inference procedures remain the same, except that only the selected rather than the complete sample is used.

[67] It might be of relevance to the empirical researcher that the variance estimates reported by Stata's *qreg* command are based on the *iid* assumption (Rogers, 1993). Furthermore, Koenker and Hallock (2000, p. 16) comment on an "unfortunate choice of bandwidth" and state that "a consequence of the undersmoothing implied by the Stata rule is that the resulting standard errors are frequently considerably smaller than would be obtained

The remaining conditions discussed in Powell (1984, pp. 310-12) and Chernozhukov and Hong (2002, p. 876), most notably independence of observations and the conditional quantile restriction (6), ensure consistency and asymptotic normality. For finite $l \geq 1$,

$$\sqrt{N}(\hat{\beta}_l(\tau) - \beta(\tau)) \xrightarrow{d} \mathcal{N}(0, M_0(\tau)), \tag{25}$$

where

$$M_0(\tau) = L_0(\tau)^{-1}\Lambda_0(\tau)L_0(\tau)^{-1} \tag{26}$$

$$L_\psi(\tau) = \mathbb{E}[N^{-1}\sum_{i=1}^{N} f_{i,\tau}(0|X_i)\mathbf{1}(h(X_i, C_i) > \tau + \psi)X_i X_i'] \tag{27}$$

$$\Lambda_0(\tau) = \tau(1-\tau)\mathbb{E}[N^{-1}\sum_{i=1}^{N} \mathbf{1}(h(X_i, C_i) > \tau)X_i X_i'] \tag{28}$$

and $f_{i,\tau}(\cdot\,|X_i)$ is the density function of the "error" $u_i(\tau) = Y_i^* - X_i'\beta(\tau)$.

The definition of the matrix $L_\psi(\tau)$ may inspire interest in the impact of the choice of k and δ_n on the efficiency of the estimator: The above variance matrix will be attained for any choice of initial estimator $\hat{\beta}_0(\tau)$ as long as $\delta_n \searrow 0$ and $|\hat{\beta}_0(\tau) - \beta(\tau)|/\delta_n \xrightarrow{p} 0$.

In contrast to the *iid* case, the variance-covariance matrix now depends on the error density at the τ-quantile, which may vary across the heteroskedastic observations. Similar to the fate of OLS in a heteroskedastic setting, the proposed estimator is not efficient. Newey and Powell (1990) derive the semiparametric efficiency bound and propose an optimally weighted estimator to attain it with weights proportional to the density at zero, $f_{i,\tau}(0|X_i)$. However, computation of the weights requires estimation of the error density; furthermore, their procedure relies on k-nearest neighbor methods. Therefore, a choice of bandwidth and k must be made, affecting not only covariance but also parameter estimates. In a travel expenditure application, Melenberg and van Soest (1996) report some degree of sensitivity to these choices mostly for standard errors and inference and less so for coefficient estimates. In addition, an underlying sample-splitting procedure may lead to inefficient estimation of the score function in practice.

The covariance matrix formula in (25) can be consistently estimated from the data, as discussed in, e.g., Powell (1984, pp. 312-14) and Powell (1986a, p. 151) for the censored case and Weiss (1991, pp. 53-54) for uncensored data. A particular challenge lies in the estimation

with a more conventional bandwidth selection rule."

of the error density at zero, $f_{i,\tau}(0|X_i)$; Powell (1984) proposes a one-sided[68] kernel estimator requiring the choice of a bandwidth. In his simulation study, Paarsch finds that the accuracy of the formula-based estimators very much depends on the degree of censoring, dramatically underestimating the variance at 50% censoring, but performing quite well at 25%.

Due to the difficulty as well as reliability issues of the above approach and the advent of a combination of cheap computation time and new methods such as the three-step CQR estimator, the bootstrap has seen a great deal of proliferation as an easy-to-compute alternative (e.g., Buchinsky, 1994; Chamberlain, 1994; Chay and Honoré, 1998). Of the covariance estimation approaches discussed in Buchinsky (1995), only the general kernel and the design matrix bootstrap (DMB)[69] are consistent for heteroskedastic data; DMB is found to clearly outperform the kernel in simulations. For a similar study, Fitzenberger (1997b) reports somewhat too conservative standard errors for the bootstrap, particularly for heavy-tailed distributions. The DMB approach carries out a nonparametric bootstrap by resampling (with replacement) the entire vector (Y_i, X_i, C_i) — in other words, drawing from the empirical joint distribution function — to yield B samples of size N. The asymptotic validity of DMB under fairly arbitrary heteroskedasticity (but independence) follows as a special case of the heteroskedasticity and autocorrelation consistent (HAC) moving block bootstrap (MBB) proposed by Fitzenberger (1998).[70]

There are two distinct approaches to constructing confidence intervals (Buchinsky, 1995, p. 308): One option is to compute an estimate $\hat{M}_0(\tau)$ as the variance-covariance matrix of the bootstrap estimates $\hat{\beta}^1(\tau), \ldots, \hat{\beta}^B(\tau)$ and derive confidence intervals from the underlying (asymptotic) normality.[71] Alternatively, the direct percentile method constructs a $1-\xi$ confidence interval from the $B\xi/2$ and $B(1-\xi)/2$ order statistics of the bootstrap estimates. Instead of resampling a complete set of N observations, an m-out-of-N ($m \ll N$) bootstrap can be performed; in this case, the variances need to be appropriately rescaled (Buchinsky, 1994; Abrevaya, 2001).

One may wonder how much more restrictive three-step CQR is in comparison with Powell's

[68] In particular, for left-censoring, Powell (1984) uses only positive uncensored residuals (negative for right-censoring). Buchinsky (1995) runs a Monte Carlo study and finds that a two-sided estimator performs better. However, it appears that in his one-sided estimator, he uses only positive residuals for a right-censored application, which seems flawed. Since these positive residuals may themselves be right-censored, this may lead to an overestimate of the density at zero and a corresponding underestimate of the standard errors, particularly at high quantiles. The results in Buchinsky (1995, table 15), however, do not support this. On the other hand, this kernel definition would not be the only typographical error in the paper.

[69] For applications, see, e.g., Buchinsky (1994) and Chamberlain (1994).

[70] Alternative bootstrap approaches include resampling based on the subgradient condition (Bilias et al., 2000) and a weighted bootstrap (Chernozhukov et al., 2011b).

[71] See also Chamberlain (1994, pp. 191-92).

canonical CQR model. For instance, does the parametric first-step classification model invoke additional assumptions which are at odds with, e.g., the heteroskedasticity in our wage data? Since the initial model only serves to provide an envelope as an incorrect lower bound of the true propensity score, no such assumptions are required; in particular, neither a particular distribution nor conditional homoskedasticity nor a location-scale submodel are imposed.

In practice, Chernozhukov and Hong report good finite-sample properties in a variety of situations. Their results are rather insensitive to the choice of probability model in step 1; we therefore use a probit model in all our applications. For their applications, repetition of step 3 did not typically bring about large improvements in terms of the Powell objective function (20), if at all. In fact, it can be useful to compare the function value of step 2 and step 3 and back up where beneficial.

Additional iterations in step 4 are akin to Buchinsky's ILPA method. Their usefulness at least partially depends on the dimensionality of the regression. When estimating a human capital model on the IABS data including only the standard controls discussed in section 4.1 — years of education, a cubic experience term, and a quadratic tenure term, but no occupation dummies — step 4 generally does not bring about large improvements in terms of the Powell objective function (20). However, in our main specification, the estimation results of which are discussed just below, we include dummies to account for the effects of 130 different occupations. In this case, step 4 still turns out to yield substantial improvements.[72] Additional iterations generally lead to only quite small or minuscule improvements, or even an increase in the objective function. We allow for three additional iterations in step 4, and select the estimates corresponding to the lowest value of the Powell function.

4.4 Three-step CQR of Mincer equation on IABS data

We estimate the Mincer equation discussed in section 4.1 on a cross section of the IABS data for 2004. Since the wage variable is top-coded, we estimate the model using (1) a Tobit ML approach, and (2)-(6) the three-step CQR estimator[73] in the above configuration at the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th} percentile (table 2). Using the direct percentile bootstrap method, we construct 95% confidence intervals, which we use to infer standard errors by appealing to

[72] The attentive reader may have correctly conjectured that these improvements are generally largest for the quantiles most affected by censoring. This is due to the fact that their sample selection is most sensitive to changes in the estimates.

[73] We use an extended version of the user-written Stata command *cqiv* (Chernozhukov et al., 2011a).

asymptotic normality.

Controlling for occupation, most of the variables have a fairly constant effect across quantiles. Heterogenous returns to education, as assumed by the accounting identity model, are likely mainly realized through the channel of occupational sorting. While we observe the familiar concave effect of experience, there is no clear-cut evidence that variation follows any particular experience pattern. Interestingly, variation is decreasing in tenure; we would instead expect that as employers acquire more knowledge about their workers, variation increases. However, the difference is rather small.

Under normality and homoskedasticity, the Tobit estimates should be very close to the CQR estimates at the 50^{th} percentile. However, most of the estimates in column (1) and (4) are several standard deviations apart, suggesting that these assumptions are not valid and the Tobit estimates biased. We further perform OPG conditional moment tests as described in section 4.2; the resulting test statistics follow $\chi^2(r)$ distributions with $r = 2$ and $r = 273$ for normality and homoskedasticity, respectively. The test statistic for normality, 4,099.3, is far larger than both the theoretical 1% critical value of 9.2 and a bootstrapped one of 7.1. The same goes for homoskedasticity with a test statistic of 14,768.0 against 1% critical values of 290.6 and 183.7. In contrast to the findings of Drukker (2002), the OPG tests here have greater size using bootstrapped rather than theoretical asymptotical values.

5 Measurement of dispersion

The preceding section has made the case for semiparametric methods to estimate censored models with considerable flexibility. What remains to be found is a reliable approach to estimating the wage dispersion within occupations. In the literature, statistics are typically computed based on the residuals from a first-stage wage regression; such approaches are discussed in section 5.1. Section 5.2 presents a new method which fully appreciates the richness of QR and controls for the dispersion effect of covariates, with simple computation properties even in the case of censoring.

5.1 Residual-based statistics

In order to infer the riskiness of occupations, Bonin et al. (2007) estimate the occupation-level standard deviation of the residuals from a Mincerian regression. This measure has frequently

been used in studies on compensation for wage risk (McGoldrick, 1995; Hartog et al., 2003; Hartog and Vijverberg, 2007; Berkhout et al., 2010).[74]

Once we move from a linear to a quantile model, residuals can no longer easily be viewed as approximating an error process. Therefore, interpretation of statistics based on residuals requires greater care. In particular, standard deviation and variance, which are used to parameterize many kinds of linear models, do not have such a simple sample interpretation anymore.[75] The grand exception is given by median regression in the case of symmetric error distributions, where, *a priori*, QR and linear methods coincide.

One may argue that, particularly with a reduced-form approach, such structural considerations are irrelevant as long as the residual-based statistic is a good proxy of the type of dispersion we intend to measure. Even then, the top-coding of the wage variable in our data leaves us with the challenge of dealing with censored residuals. Cizek (2008) assumes normality and proposes to estimate the standard deviation using the median of negative uncensored residuals:[76]

$$\hat{\sigma}_n = 1.4826 \cdot \text{med}\{\hat{u}_i : Y_i < X_i'\hat{\beta} < C_i\} \tag{29}$$

In doing so, Cizek (2008) uses the well-known relation between standard deviation and median absolute deviation (MAD) under normality[77] as well as symmetry. Rousseeuw and Croux (1993) discuss properties of MAD-based estimators and present two alternative estimators.

An alternative approach may be based on the Kaplan-Meier product limit estimator (Kaplan and Meier, 1958), which is commonly used in the analysis of duration or survival data and gives the nonparametric maximum likelihood estimate of the survival function. In an application to residuals, the survival function may be interpreted as the (conditional) cumulative distribution function (cdf) and the corresponding Kaplan-Meier estimator as the empirical distribution function (edf). Chay and Honoré (1998) and Chay and Powell (2001) use it to estimate the error distribution on residuals from CLAD wage regressions and find long tails compared to a normal distribution. However, this approach is problematic under covariate-dependent heteroskedas-

[74] See also the overview paper by Hartog (2009) on risk-augmented Mincer equations.

[75] For certain distributions, such as the Cauchy distribution, the variance does not even exist. In that particular case, neither does the mean, such that the necessary assumptions for OLS are not valid, which may strengthen the case for methods such as QR (see Paarsch (1984) for Monte Carlo evidence on parametric and semiparametric methods). Smith (1973) provides a theoretical as well as practical discussion of the behavior of OLS estimates for Cauchy errors.

[76] Actually, Cizek (2008) assumes censoring from below, and he therefore uses positive uncensored residuals. Since any censoring problem can be transformed into the classical setting of left-censoring at zero, this goes without any loss of generality. Cf. the kernel density estimator in Powell (1984).

[77] $\sigma = 1/\Phi^{-1}(\frac{3}{4}) \cdot \text{MAD} \approx 1.4826 \cdot \text{MAD}$

ticity, since the censoring point will be informative for the error distribution,[78] resulting in inconsistency of the Kaplan-Meier estimator. The Buckley-James estimator for censored models (Buckley and James, 1979) applies the Kaplan-Meier estimator to the error distribution using the residuals and their censoring points.[79] Naturally, it is not robust to heteroskedasticity (e.g., Powell, 1984; Honoré et al., 2002); while the censoring points may depend on the regressors, the errors may not, such that the latent dependent variable is independent of the censoring.

Just as conditional expectations may only insufficiently characterize the wage distribution in its totality, statistics such as the standard deviation may be too simplistic to describe the distribution of errors. Hartog and Vijverberg (2007) additionally consider the interquartile range of residuals, i.e., the distance between the 25^{th} and the 75^{th} percentile. However, this mostly appears to serve robustness and does little to illustrate possible asymmetry or shape differences in the error distributions in different occupations. Furthermore, one should not mistake the analysis of residual quantiles from OLS or LAD regression as equivalent to QR methods; this is illustrated by Bassett and Koenker (1982, pp. 413-14).

5.2 Differences of quantile coefficients

This section presents a new method for the estimation of dispersion which is not based on residuals, but rather on the difference of coefficient estimates at particular quantiles. As such, it is in the spirit of the heteroskedasticity test of Koenker and Bassett (1982), which carries out a Wald test on the differences of coefficient estimates at different quantiles. Specifically, we first estimate the entire model by (C)QR at different quantiles, such as the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th} percentile, including dummy variables for the groups which are to be compared; in our application (section 4.4), for instance, we include dummies for all occupations. We then consider the coefficient estimates for these dummies at two particular quantiles, such as the 10^{th} and 50^{th} percentile, and compare their differences across occupations. The following sections will discuss this approach in the context of a location model (5.2.1), a location-scale model (5.2.2), and a more general model (5.2.3).

Our approach is not only computationally simple, but it also controls for the dispersion effect of covariates, and thereby filters out the (possibly) heteroskedastic effect of, e.g., education and tenure. This is particularly important in our case, since we match two different samples: While

[78] Specifically, since the upper wage limit is constant in our sample, the censoring point of the residuals is a function of the covariates.

[79] For a brief exposition and discussion, see Powell (1994, pp. 2504-05).

Bonin et al. carry out all estimation on the SOEP data and can therefore control for these effects in their second-stage regression, this would only work here if our samples from the IABS and SOEP data perfectly matched in terms of their distribution.[80]

5.2.1 Location model

Suppose initially that the random variable $Y_i = (Y_1, \ldots, Y_n)$ is generated by a linear model with *iid* errors u_i:

$$Y_i = X_i'\beta + u_i \Leftrightarrow \Pr(Y_i \leq y|X_i) = F(y - \sum_{j=1}^{p} X_{ij}\beta_j) \tag{30}$$

Recall our assumption that $X_{i1} = 1$, $i = 1, \ldots, n$, i.e., our regressors always include a constant. The CQF is given by

$$q_\tau(X_i) = X_i'\beta(\tau) \tag{31}$$

$$= X_i'\beta + F^{-1}(\tau) \tag{32}$$

Therefore, $\beta(\tau) = \beta + (F^{-1}(\tau), 0, \ldots, 0)$, and the quantiles affect the process only by a shift of location.

5.2.2 Location-scale model

In a location-scale model, Y_i is additionally dependent on X_i through a re-scaling of variances (the disturbances u_i remain *iid*):

$$Y_i = X_i'\beta + \sigma(X_i)u_i \tag{33}$$

Under this form of systematic heteroskedasticity, conditional quantiles change accordingly:

$$q_\tau(X_i) = X_i'\beta + \sigma(X_i)F^{-1}(\tau) \tag{34}$$

Suppose $\sigma(\cdot)$ is linear with $\sigma(X_i) = 1 + X_i'\zeta$.[81] Then, $\beta(\tau) = \beta + F^{-1}(\tau)\zeta + (F^{-1}(\tau), 0, \ldots, 0)$,

[80] Both samples are representative of the German population and workforce, respectively, in the way described, and we can harmonize them by, e.g., excluding civil servants in both. However, due to small sample size, particularly in some occupation cells, the SOEP may not be quite as useful to control for the covariates' effect on wage dispersion.

[81] Cf. the linear scale model of heteroskedasticity in Koenker and Bassett (1982, p. 45, (2.12)).

and the quantiles have both a location and a scale effect. To strike the link to our problem, suppose that in addition to a vector $\tilde{X}_i \in \mathbb{R}^p$ of constant and $p-1$ covariates, X_i contains a vector $\dot{X}_i \in \mathbb{R}^{m-1}$ of dummies for m occupation groups such that $X_i = (\tilde{X}_i, \dot{X}_i)'$; in this case, $\gamma_{p+k-1} = \dot{\sigma}_k$ denotes the occupation-specific scale effect of occupation group $k \in \{2, \ldots, m\}$.[82] Then, for consistent estimators $\hat{\beta}(\tau_1)$ and $\hat{\beta}(\tau_2)$, $\hat{\beta}_{k+p-1}(\tau_1) - \hat{\beta}_{k+p-1}(\tau_2) \xrightarrow{p} \dot{\sigma}_k[F^{-1}(\tau_1) - F^{-1}(\tau_2)]$ by the continuous mapping theorem (CMT). Hence, we can consistently estimate $\dot{\sigma}_k$ up to scale (no pun intended), and after imposing normalization or a parametric specification, its precise value can be obtained.

The model described above is comparable to the normal-location models discussed by Chamberlain (1994, p. 186). Instead of our single model, in which occupations affect location and scale only in an additive way, Chamberlain specifies two separate models for the union and non-union sector. He shows how a difference in variance between the two models affects the estimated union wage effect across the quantiles. However, Chamberlain considers this normal-location model inadequate for characterizing the conditional wage distribution. In particular, it implies constant covariate slopes across the quantiles; this is at odds with the quantile patterns of industry wage effects Chamberlain finds. Moreover, and closest to our analysis, he presents differential patterns across industries and relates these to industry-specific residual dispersion. Close inspection of the industry coefficients in Chamberlain (1994, table 5.4) at different quantiles reveals that even the location-scale model may be too restrictive. This point is explored and elaborated on in the following.

5.2.3 General (separable) dispersion model

The above location-scale model allows some flexibility with respect to heteroskedasticity, but it still imposes a great deal of homogeneity on the shape of the error distributions. In particular, by the CMT, $\frac{\hat{\beta}_{p+k-1}(\tau_1) - \hat{\beta}_{p+k-1}(\tau_2)}{\hat{\beta}_{p+l-1}(\tau_1) - \hat{\beta}_{p+l-1}(\tau_2)} \xrightarrow{p} \frac{\dot{\sigma}_k}{\dot{\sigma}_l}$ for any $\tau_1, \tau_2 \in (0,1)$ and $k, l \in \{2, \ldots, m\}$.[83] For instance, differential asymmetry and tail behavior are precluded.

Consider therefore a more general linear model in which each occupation $k \in \{1, \ldots, m\}$ has its own conditional distribution function, such that $u_i \sim F_k(\cdot\,; X_i)$ if individual i works in occupation k. Suppose further that for each k and at every quantile τ, $F_k^{-1}(\cdot\,; X_i)$ is linearly separable into a general component $\xi(\tilde{X}_i; \tau)$ linear in $\tilde{X}_i$ and an occupation-specific component

[82] We omit occupation 1 and assume that its effect is subsumed in the general constant.
[83] Cf. Koenker and Xiao (2002a, (5.1)).

$\eta_k(\tau) = \dot{X}_i{}'\eta(\tau)$.[84] The CQF is then given by:

$$q_\tau(X_i) = X_i'\beta + \xi(\tilde{X}_i;\tau) + \eta_k(\tau) \tag{35}$$

$$= X_i'\beta + \tilde{X}_i{}'\lambda(\tau) + \dot{X}_i{}'\eta(\tau) \tag{36}$$

As a result, $\beta(\tau) = \beta + (\lambda(\tau), \eta(\tau))$, and hence, for consistent estimators $\hat{\beta}(\tau_1)$ and $\hat{\beta}(\tau_2)$, $\hat{\beta}(\tau_1) - \hat{\beta}(\tau_2) \xrightarrow{p} (\lambda(\tau_1) - \lambda(\tau_2), \eta(\tau_1) - \eta(\tau_2))$. In this way, we can measure the effect of our set of occupation groups on all different quantiles while controlling for the dispersion effect of the additional covariates. This gives us a measure of the conditional dispersion effect of each occupation.

While the linear specification of the conditional quantiles may appear restrictive, a linear quantile model is frequently only intended as a reduced-form approximation, such as for the minimum distance (MD) estimators in Buchinsky (1994, p. 409) and Chamberlain (1994, p. 181).[85] Also, the classical QR-based heteroskedasticity test of Koenker and Bassett (1982) uses a linear approximation to the variance structure; the authors discuss how, for example, a model with multiplicative heteroskedasticity can be approximated by a simple linear expansion. As an alternative, Chamberlain (1994, pp. 187-91) presents a Box-Cox procedure for QR, which allows the researcher to generalize the functional form of the linear model and relies heavily on the equivariance property.

More formally, Chamberlain (1994, p. 181) recognized that the QR estimator provides a linear approximation to the CQF, albeit of a less "transparent" nature than in the OLS and MD case. Angrist et al. (2006) take up this issue and show that QR minimizes a weighted mean-squared error loss function for specification error, implicitly providing a weighted MD approximation to the true nonlinear CQF. Applying the framework to wage regressions with a focus on the education variable, they find QR to provide a useful approximation to the conditional wage distribution.

In working with log income, we acknowledge the likely nonlinear nature of income. Not only

[84] This does not state that $\eta_k(\tau)$ is indeed the τ-quantile of the occupation effect (cf. the concept of comonotonicity (Koenker, 2005, p. 60)); rather, it is the occupation effect at the τ-quantile. Therefore, $\eta_k(\tau)$ need not be increasing; rather, an occupation l may experience lower variance than the reference occupation, in which case $\eta_l(\tau)$ is decreasing. In fact, the estimated quantile functions will not generally be monotonic in τ at all sample points ("quantile crossing"), although at the mean (Koenker, 2005, pp. 55-59); Chernozhukov et al. (2010) provide a correction method. The linear specification of our model (possibly as an approximation) is not directly tied to this issue and will be discussed in more detail below.

[85] See also the discussion in Fitzenberger (1997b, p. 422) and Abrevaya and Dahl (2008, pp. 381-82, 394).

do we allow for a possibly more natural interpretation of the slopes as relative changes,[86] but we also implicitly assume an interaction effect of the covariates with respect to the level of income. Due to equivariance and distribution freeness of QR, we enjoy somewhat greater flexibility with respect to the specification of the dependent variable, though this is by no means a trivial issue. Conversely, while the use of log income in OLS estimation is frequently motivated by the observation that wages are approximately log-normally distributed, this is of less importance here.

To evaluate the practical merits of the approach, we carry out a set of Monte Carlo simulations (section B in the appendix). We find that our method is particularly effective, when compared to conventional approaches, for the estimation of dispersion in the presence of interaction effects in variance.

6 Results on cross-sectional risk

This section will empirically test for a relationship between the risk attitudes of individuals in the SOEP and the earnings dispersion in their respective occupations. In section 6.1, earnings dispersion is estimated on the SOEP and then matched to all individuals on an occupation basis, while in section 6.2, the same risk profiles are estimated on the IABS data and matched to individuals in the SOEP as described in section 3.3. In both cases, we follow the approach of Bonin et al. (2007) and use simple OLS regressions positing the earnings risk statistic of interest as the dependent variable and risk attitudes as well as a set of control variables as regressors.

6.1 Replication of Bonin et al. (2007)

We initially reconsider the (uncensored) data of Bonin et al. (2007) from the 2004 wave of the SOEP using the approach described in section 5.2. Recall that in contrast to our main specification, we include individuals from East Germany as well as tenured civil servants. As a measure of dispersion, we use the difference between the occupation coefficients (at the ISCO 2-digit level) at the 90th and 10th percentile based on QR using the same specification as in Bonin et al. (2007, table 1, (2)).[87] The results are reported in table 3. The coefficient estimates

[86] This also facilitates comparison of the difference estimates. If a higher wage level causes greater variation, as appears plausible, it may be more sensible to work with relative changes across quantiles in order to avoid any spurious relation which might arise if more risk tolerant individuals work in better paid professions due to any of the points discussed before.

[87] This difference will hereafter be referred to as the "10-90 spread," not to be confused with the *.90-.10 spread* of Buchinsky (1994, p. 444), which is defined as the difference between the conditional quantiles from group-based

for risk attitude are significant (in the expected direction) at the 5% level in every specification and allowing for clustering at the occupation level, which substantially increases standard errors. Consistent with a reduction in measurement error and attenuation bias, the coefficient estimates on average risk attitude are larger than those based on the single measure.

To get an idea of the size of the effect in comparison to previous work, it is necessary to assume an underlying distribution; under normality, we estimate a one-point increase in risk attitude to result in an increase of .0007 or .001 standard deviations in earnings dispersion, based on specifications (2) and (5).[88] These effects are similar to those reported in Bonin et al., as are the R-squared statistics. The current approach can be regarded as more robust, as it filters out the effects of the covariates at the different quantiles of the human capital model, while additionally controlling for these characteristics in the above risk regression. The coefficient estimates of the control variables are difficult to interpret; they may suggest specification error in the wage regressions, but can also reflect the desire for occupational risk of the different groups. Occupational earnings dispersion and the level of earnings are positively related, as may have been expected. Height also enters positively, which is in line with results that taller people are more risk seeking (Dohmen et al., 2011b). Married individuals work in lower-risk occupations, possibly due to greater household size and concomitant financial commitments, while more educated individuals work in occupations with greater earnings risk. The positive coefficient estimate on public sector employment is difficult to interpret; certain occupations may coexist in the public and private sector, with selection out of the public sector (to the bottom and to the top) leading to comparatively high earnings dispersion.

The 10-90 spread is useful to characterize the overall dispersion; however, if the location-scale model assumption is not valid, it will be fruitful to take a more disaggregate look as well. The location-scale hypothesis can be scrutinized using the Khmaladze test proposed in Koenker and Xiao (2002a), which computes a distribution-free test statistic on the QR process using a Khmaladze transformation. This test statistic needs to be compared to tabulated critical values. Since these are only available for a limited number of covariates, we switch to the broadest ISCO 1-digit level; in the Mincer specification, we use the general risk response of 2004 as a measure of risk tolerance. Using a .05-spaced grid from the 10th to the 90th percentile, we compute a test statistic of 25.87,[89] which is larger than the 1% critical value of 20.21 (Koenker and Xiao,

estimation.

[88] Cf. Buchinsky (1994, fn. 35).

[89] We use the Khmaladze test included in the package *quantreg* for R.

2002b, p. 5, $p = 18, \epsilon = .1$) and thus compels us to reject the location-scale hypothesis at the 1% level.

We therefore turn our attention to a set of different percentiles, and run the same regressions for the differences between the estimates at the 50th percentile and the 10th, 25th, 75th, and 90th percentile. The results, again allowing for clustering, are reported in table 4. The relationship between risk attitudes and earnings dispersion appears clearer at the lower end of the distribution, reflected by the greater significance of estimates in the specifications corresponding to the 10-50 and 25-50 spreads. This would suggest that downside risk is more relevant.

Using the KldB 88 classification of the IABS data, the results carry through for the 10-90 spread, while we find no significant effect for the various median differences. However, the analysis based on SOEP wage data is rather crude, due to the small cell sizes even at the rather coarse 2-digit level. The IABS data, on the other hand, offer much greater cell sizes even at a finer level of coding, which appears crucial to the precise computation of quantiles.

6.2 IABS risk profiles

6.2.1 Main results

In this section, we no longer estimate the risk profiles of occupations on the SOEP, but on the separate IABS data, in particular a cross section for 2004. They hence follow from the regressions presented in section 4.4.

Again, we depart from a specification using the 10-90 spread as a general measure of dispersion; results are reported in table 5. Instead of log monthly income, we include the median wage of the occupation as a control,[90] since it is more plausibly known by the individual and considered when choosing an occupation; furthermore, its coefficient may reflect a relationship between level and dispersion.

We estimate a positive effect of risk tolerance in all specifications, significant at either the 5% or the 10% level; as before, the effect is larger for average risk tolerance, likely due to greater measurement precision.[91] The estimated effect size is larger than in the above replication with

[90] For a few occupations (electrical engineers, managers/executives, doctors and pharmacists), median earnings are above the censoring threshold, and hence, median wage is censored. Omitting these occupations does not affect the results for overall dispersion and the upper regions, but slightly reduces the significance of the estimates at the lower quantiles.

[91] The number of available responses to the general risk question differs by individual between one and four and concurrently the precision with which we estimate the latent true risk preferences. While not theoretically necessary, one may find it desirable to account for this by applying weights. Denote the available measurements for true preferences α_i by $\lambda_{i,j}$ for $i \in \{1, \ldots, N\}$ and $j \in \{1, \ldots, k_i\}$, and assume that $\lambda_{i,j}|\alpha_i \overset{\text{iid}}{\sim} (\alpha_i, \sigma_\lambda^2)$. Then,

an increase of .0009 and .0018 standard deviations per point increase on the 11-point risk scale for the single-year and average measure of risk tolerance, respectively (again, assuming normality for purely expositional purposes). The R-squared statistic on the other hand is notably smaller; it is conceivable that this is a statistical artefact caused by the fact that above, both stages of the estimation were carried out on the same data (SOEP). Of the control variables, only marital status, education, and median wage are significant. As above, married individuals work in less risky occupations.

Surprisingly, median wage now enters negatively; it seems unlikely that this is entirely due to wage compression in the top regions of occupations with high median wage. Instead, observations in high-wage occupations are naturally more likely to be censored. If censoring leads to an underestimation of dispersion, this will be picked up by the coefficient on median wage. The effect of median wage is clearly the most pronounced at the 90^{th} percentile (table 7), which is affected most by censoring; in fact, the unconditional 90^{th} percentile is censored (table 1). As long as there is sufficient within-occupation heterogeneity that each of them will contain a number of individuals with uncensored conditional 90^{th} percentile, this will not be a problem given correct specification of the model. Since median wage has a standard deviation of around .3, the magnitude of the distortion is not that large in principle; however, it is a lot larger than the effect of risk attitudes. Reassuringly, these are very similar when considering the 25-75 spread instead (table 6).

Turning our attention to individual median differences, the results in table 7 do not show any difference between the lower and the upper part of the wage distribution; all quantiles are similarly correlated with risk attitudes. Also here, we find a larger coefficient estimate on the average measure of risk tolerance.

With respect to the three-step CQR method, the significance of the second-stage estimates just presented is sensitive to the precision of our wage regressions. As discussed at the end of section 4.3.2, substantial improvements are possible by going beyond the standard number of steps; specifically, we allowed for an additional three iterations to reduce measurement error in the left-hand side variable of our second-stage regressions. However, rather than decrease the standard errors of our second-stage estimates, this mostly increased the magnitude of the coefficient estimates; the prior measurement error may not have been exogenous due to covariate-

$\sum_{j=1}^{k_i} \lambda_{i,j}/k_i \sim (\alpha_i, \sigma_\lambda^2/k_i)$, and one option is to use weights proportional to $\sqrt{k_i}$. However, doing so has very little effect on our estimates.

dependent censoring probabilities. The results using additional iterations also turned out to be more robust to alternative specifications, as discussed below.

6.2.2 Domain-specific risk attitudes

In addition to general risk attitudes, the SOEP waves of 2004 and 2009 contained similarly phrased risk questions for specific domains, among them career.[92] We find the general risk question preferable for our analysis due to its somewhat closer links to established concepts in economics and psychology on the one hand and the possibly greater danger of endogeneity and more suggestive nature[93] of the career risk question on the other hand. Still, we would like to see how the career risk question performs in the context considered.

Results for the 10-90 spread and median differences are reported in tables 8 and 9, respectively. In most cases, the significance of the 2004 career risk question is more pronounced than that of the general risk question. In contrast to average general risk attitudes, average career risk attitudes actually explain slightly less variation than the 2004 response. This is surprising given the results for general risk, and since we would expect the second response to be the most incrementally useful one. The reason may be a large drop in risk tolerance in 2009, likely due to the financial and economic disruption at that time, which was mostly a rank-order preserving location shift affecting the entire distribution.[94] Due to panel attrition and other factors, the 2009 response is only available for around two thirds of the individuals; for the rest, average risk tolerance is equal to 2004 risk tolerance. However, if we only use individuals with responses to both the 2004 and 2009 question, the effect is still slightly lower than for the 2004 response alone. At around .0011 standard deviations per point increase on the scale, the effect is roughly that measured for average general risk tolerance. Constructing the latter average from 2004 and 2009 responses only, we find a much smaller effect than using all four available responses, but it is nonetheless greater than that of the 2004 general risk response alone. Possibly, measurement error in the career risk question is more persistent than in the general one.

[92] The exact wording (translated from German) is as follows: "People can behave differently in different situations. How would you rate your willingness to take risks in the following areas?"

[93] When asked for their career risk preferences, individuals might use a sort of revealed preference approach and think about the riskiness of their present occupation. This type of framing seems less relevant for the general risk question.

[94] Therefore, the stability assumption is not invalidated, at least in the ordinal sense.

6.2.3 Specification of the education variable

There are two dimensions along which we may vary the education variable in our estimation, one with respect to model specification, the other with respect to the data. Particularly in the German context, schooling may have a nonlinear effect on earnings. To model this, we substitute years of education by a set of dummy variables indicating the completion of different schooling levels. However, this change has hardly any effect on the results presented just above.

Our main analysis uses a schooling variable constructed by an algorithm to correct for misreporting. It turns out that this correction has an effect on the significance of our estimates, as without it, coefficient estimates are noticeably smaller. However, it appears more plausible for misreporting to result in artificial variation in the human capital model than for the correction to induce any effect driving up the coefficient estimates on risk tolerance in an entirely separate data set. In any event, the results for the 25-75 spread reported in table 10 suggest an effect even in case the original education reports are (more) correct.

7 Discussion

7.1 Causality

Our analysis requires particular assumptions for a causal sorting interpretation, most prominently, exogeneity of occupational risk and risk attitudes. Even then, it is still possible that the driving force is not individuals sorting into occupations according to risk, but employers catering to the risk preferences of their (potential) employees in wage bargaining. Section 7.3 provides some results on the temporal stability of occupational risk profiles, which suggest that there has been relatively little adaptation.

We could extend our analysis by reviewing the evidence for different groups. High-ability individuals, for instance, may have a richer menu of occupations to select from along with a greater aptness at evaluating information on the riskiness of occupations and their own risk preferences. East Germans may be less efficiently sorted due to the absence of earnings risk at the time of sorting (Fuchs-Schündeln and Schündeln, 2005). Bonin et al. (2007) find a hump-shaped age-correlation profile for risk preferences and occupational risk peaking at mid-career; Gibbons and Murphy (1992) show that in a multiperiod model with risk aversion and imperfect knowledge about workers' productivity, firms will insure against bad realizations by offering wages independent of ability at early career stages, later introducing performance-based pay to

stimulate effort. Any analysis on subgroups is unfortunately troubled by small sample sizes.

Instead, it may be useful to consider the role of individuals' information. In our second-stage regression, we assume that individuals are aware of the existence of heterogeneity, in line with the accounting identity model of Mincer (1974), but have no private information on their ability aside from that revealed in their choice of education. Unlike the econometrician, however, to whom heterogeneity is simply unobserved, individuals may hold private information on their abilities, but also on softer factors such as their motivation. Focusing on schooling choice, Cunha and Heckman (2007) survey a framework of recently developed methods to separate ex ante and ex post returns to schooling, and find that roughly half of the earnings variation in their NLSY sample is known to agents in advance. However, it appears daunting to extend their approach, based on the Roy model (Roy, 1951), to occupational sorting, requiring counterfactual estimation of a large number of earnings streams.

Nonetheless, we cannot yet rule out an effect driven by cognitive abilities rather than risk preferences. As discussed in the literature section, there is evidence for a negative relationship between risk aversion and cognitive abilities; at the same time, dispersion is particularly attractive for high-ability individuals. In our general framework, there are two distinct approaches to account for private information on abilities, with different economic as well as econometric ramifications. The first one is based on our models and estimation techniques hitherto and assumes that for each individual, only a certain part of the conditional wage distribution is relevant; for instance, high-ability types may only consider the 50th to 90th percentile when making their decision, and therefore only consider the related risk.[95] Alternatively, we may assume that people have perfect knowledge of their individual effects within each occupation, and that all risk borne arises from frictions or intertemporal variation. This is obviously suggestive of an approach incorporating fixed or other individual-level effects, which is pursued in section 7.2.

7.2 (Un)observed heterogeneity

Honoré (1992) introduces QR-based estimation techniques for censored and truncated data consistent for a finite longitudinal dimension T, but they are computationally intensive and do not appear feasible for our application;[96] furthermore, Honoré only discusses estimation of the

[95] The mindful reader will note that for a location-scale model, unless the range of relevant quantiles differs by occupation — which is not implausible, but very hard to identify without measures of skills or ability and a large number of observations relative to the number of occupations (the SOEP fulfills the former criterion, and the IABS the latter; neither does both) — the analysis will not differ from that in section 5.2.

[96] A useful discussion of this and other methods is given in Honoré and Kyriazidou (2000) and Hsiao (2003, pp.

median, but of none of the noncentral quantiles. Honoré (1993) and Honoré and Hu (2004) consider dynamic panel data models; however, their methods are only applicable for censoring arising from corner solutions, but not top-coding of wages.

One of the fundamental issues in our analysis is that censoring introduces nonlinearity. However, if comonotonicity fails, quantiles may experience nonlinearity even if the data are not censored. Hence, since conditional quantiles are not linear operators, conventional differencing strategies will fail to eliminate individual effects. As Koenker and Hallock (2000, p. 19) put it: "Quantiles of convolutions of random variables are rather intractable objects, and preliminary differencing strategies familiar from Gaussian models have sometimes unanticipated effects." The literature on panel methods for uncensored data includes contributions by Koenker (2004) and Ponomareva (2011); while the former restricts the individual effects to location shifts, the latter allows for the fixed effects to vary over the quantiles. However, although Ponomareva discusses estimation, we are not aware of an implementation of the approach, based on a sequence of moments estimators, in either her paper or elsewhere. The approach of Koenker (2004), based on penalized fixed effects and a shrinkage parameter,[97] appears computationally problematic for large N and small T due to the nature of the objective function (Bache et al., 2011; Dahl et al., 2009). Abrevaya and Dahl (2008) present a method consistent for fixed T and allowing the individual effects to differ across quantiles. Due to its relation to the correlated random-effects model of Chamberlain (1982, 1984), it imposes a particular structure on these effects; while this need not be too restrictive, estimation might be troubled by identification issues (Canay, forthcoming).

We therefore chose to implement a simple two-step estimator for uncensored panel data devised by Canay, which views fixed effects as location shifts. In the first step, we fit a standard fixed-effects model and estimate the individual effects as the individual-level averages of the residuals from the linear projection based only on the covariates. After subtracting these, we estimate a pooled model using QR.

The specification of the individual effects requires care. One option is to estimate both individual and occupation effects; however, this would mean that individual effects are identical across occupations, which seems to be at odds with our earlier analysis on uncertainty regarding

243-59). Attempts to estimate a wage model including fixed effects using the Stata command *pantob* (Honoré, 2009) based on Honoré (1992) were to no avail. See Chay and Honoré (1998) for a successful application.

[97] See also Koenker (2005, pp. 278-81) for a brief introduction.

the wage distribution. We therefore select a specification including individual-occupation effects:

$$Y_{ikt} = X_{it}'\beta + c_k + \eta_{ik} + d_t + u_{ikt} \tag{37}$$

$$= X_{it}'\beta + v_{ik} + d_t + u_{ikt} \tag{38}$$

for individual $i \in \{1, \dots, N\}$ in occupation $k \in \{1, \dots, m\}$ at time $t \in \{1, \dots, T\}$; furthermore, $u_i \sim F_k(\cdot\,; X_{it})$ with $\mathbb{E}(u_{ikt}|X_{it}, \eta_{ik}, c_k) = 0$.

We do not make any distributional assumptions on the term η_{ik} except that it acts as a location shift with identical effect across quantiles, but this is inherent in the linear specification. In the first step, we cannot simultaneously identify both the occupation and the occupation-individual effect. Rather, we estimate and difference out the term v_{ik} as described above. In the second step, we can again include occupation dummies; differencing out their effect in the preceding step only leads to a location shift which is constant across quantiles. Performing QR at different quantiles will then provide us with statistics of the occupation-specific dispersion after accounting for individual effects.

The SOEP waves from 1991 through 2009 are used to construct a sample of full-time employed men aged 25 to 55, as in the replication of Bonin et al. (2007) in section 6.1. We deflate (uncensored) wages according to the German consumer price index (Statistisches Bundesamt Deutschland, 2011) with base year 2004. Since consistent estimation of the individual effects requires $T \to \infty$,[98] we only use individuals with at least five observations. The Mincer specification again includes a linear schooling term, a cubic and a quadratic polynomial for experience and tenure, respectively, and dummies for public sector employment and East Germany as well as occupations at the ISCO 2-digit level. Furthermore, average risk tolerance is included to account for its effect as a productive trait.

After computing the by-now familiar dispersion statistics, we again perform OLS regressions to estimate the relationship of occupational earnings risk and risk attitudes in the 2004 cross section. The results for the 10-90 spread are reported in table 11. Surprisingly, the coefficient estimate on risk preferences is negative and significant, though the size of the effect is much smaller than before. Closer inspection reveals that this is driven by "switchers", i.e., individuals who were previously in a different occupation (in the sample period); for this group, there are clear effects in all parts of the conditional wage distribution (table 12). The picture remains

[98] Canay proves that the parameters of interest are identified even for fixed $T \geq 2$.

similar if we only use the 2004 cross section for the Mincer estimation after eliminating individual effects, suggesting that cyclical rather than longitudinal variation is not the main driver of this effect.

The estimation results suggest that conditional on individual effects, more risk tolerant agents actually take on lower risk. However, we are rather hesitant about such an interpretation. The effect may rather be due to differences in occupation characteristics, specifically the capacity of the employer to observe ability and performance. In occupations which allow precise screening and evaluation, wage dispersion may be large, whereas other occupations offer only average wages, which may have a closer tie to firm or industry rather than individual performance. The enhanced information on workers' productivity might actually lead to smoother and more stable compensation paths.

7.3 Stationarity of labor markets

In section 3.1.2, we discussed the stability properties of risk attitudes. Failure of the stability assumption may bias the coefficient estimates on risk tolerance in section 6 in two directions: 1) If risk preferences are endogenous and adjust to the occupational risk experienced, the magnitude of the effect will likely be overstated, 2) if risk preferences are unstable, but respond mainly to (random) shocks unrelated to occupational risk, the effect will likely be understated due to a low correlation between measured risk tolerance and risk tolerance at the time of sorting.

Since there are two sides to risk in our analysis, risk attitudes of individuals and risk profiles of occupations, it is useful to investigate the stability, or stationarity, of the latter as well; with data spanning back to 1975, we can cover a lot of ground in this respect. A lack of stability would imply that similar to point 2), the risk profile estimated in the cross section for 2004 may not be the one relevant for the (sticky) sorting decision. The effect on our regression estimates, however, would be slightly more subtle: In contrast to its effect on regressors, measurement error in the left-hand side variable does not induce attenuation bias; rather, it tends to inflate the standard errors of regression estimates.

As an admittedly crude measure of stationarity, we extend the estimation of risk profiles in section 6.2 to the years 1979, 1984, 1989, 1994, and 1999, and compute Pearson correlation coefficients with occupations as cross-sectional unit (table 13). Working with log wages, we can ignore inflation and the introduction of the *Euro* in 2002; both only cause a location shift constant across quantiles. The correlation coefficient is decreasing in the time span considered,

but remains positive and high; it is larger than .65 for any pair of years. Results are similar for the individual median differences and using Spearman rank correlation.

This suggests that the relative wage dispersion of occupations has been rather stable in West Germany in the period from 1979 to 2004, and that the risk profiles we estimate from a cross section for 2004 are quite close to those relevant at the point of job entry for most of the individuals in the SOEP. Of course, the dispersion is to some extent a result of how occupations are coded and therefore aggregated and grouped together. Since this is constant over time, it will account for some of the correlation measured. However, the sheer number of occupations in our classification (130) makes it less credible that this is a particularly great source of the observed stability.

8 Conclusion

We present evidence for a relationship between individuals' risk attitudes and the earnings risk in their occupations, supporting earlier findings of, among others, Bonin et al. (2007). In contrast to the latter, our estimation draws on two different German data sets; while individual risk attitudes and other characteristics are taken from the Socio-Economic Panel (SOEP), we use the IAB Employment Sample (IABS), composed of administrative wage data, to estimate the risk profiles of over one hundred different occupations. Along the way, we tackle several econometric challenges: Since the IABS wage data are top-coded, we use the three-step estimation method of Chernozhukov and Hong (2002) based on semiparametric methods; due to the high dimensionality of our regressions, additional iterations are crucial for the precision of our analysis. To estimate the dispersion of wages within occupations, we estimate quantile regressions including dummies for all occupations, and compute differences of the occupation coefficient estimates at different quantiles; this elegantly circumvents the problems of residual-based statistics in the presence of censoring, while controlling for the dispersion effect of covariates. The value of the latter feature proves considerable in a simulation study.

To estimate the relationship between these dispersion measures and individual risk attitudes, simple OLS regressions are used. Consistent with regression dilution and attenuation bias, we find greater effects of risk tolerance measures averaged over responses of several years. This contributes applied evidence to the discussion surrounding the measurement of risk preferences, specifically stability and measurement error.

Our analysis yields the powerful result that a relationship of risk attitudes and employment characteristics is not just manifested in the choice between entrepreneurship and dependent employment or the public and private sector. Rather, even within the private sector, more risk tolerant individuals work in occupations with greater earnings risk. Necessary conditions for a causal sorting interpretation are briefly discussed; chiefly, we require exogeneity of risk attitudes and occupational risk.

There are several lines along which our analysis may be fruitfully extended. We have so far abstracted from unemployment risk; however, unemployment certainly represents a relevant financial risk, and may even have a psychological dimension. Furthermore, while we have exploited the panel dimension of the SOEP to eliminate individual effects, there is opportunity for further research on the relationship of risk attitudes and wage paths, particularly in light of our surprising results after accounting for unobserved heterogeneity. While the IABS data are rich in both the longitudinal and cross-sectional dimension, censoring of wages creates substantial challenges for such an analysis. Hu (2002) presents a GMM estimator and uses the estimate of the autoregressive coefficient as a statistic for earnings stability for samples of blacks and whites from the March Current Population Survey and Social Security Administration (CPS-SSA) earnings record.

Finally, our paper remains silent about the decision on education, taking it as given. There are divergent hypotheses on the influence of risk attitudes, derived from viewing education as either insurance (e.g., against unemployment) or a risky investment. The joint decision on education and occupation may offer fertile ground for future research.

References

Abrevaya, J. (2001). The effects of demographics and maternal behavior on the distribution of birth outcomes. *Empirical Economics*, 26(1):247–57.

Abrevaya, J. and Dahl, C. M. (2008). The effects of birth inputs on birthweight. *Journal of Business & Economic Statistics*, 26(4):379–97.

Amemiya, T. (1973). Regression analysis when the dependent variable is truncated normal. *Econometrica*, 41(6):997–1016.

Amemiya, T. (1984). Tobit models: a survey. *Journal of Econometrics*, 24(1-2):3–61.

Amemiya, T. (1985). *Advanced Econometrics*. Harvard University Press, Cambridge, MA.

Angrist, J., Chernozhukov, V., and Fernández-Val, I. (2006). Quantile regression under misspecification, with an application to the U.S. wage structure. *Econometrica*, 74(2):539–63.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, Princeton, NJ.

Arabmazar, A. and Schmidt, P. (1981). Further evidence on the robustness of the Tobit estimator to heteroskedasticity. *Journal of Econometrics*, 17(2):253–58.

Arabmazar, A. and Schmidt, P. (1982). An investigation of the robustness of the Tobit estimator to non-normality. *Econometrica*, 50(4):1055–63.

Bache, S. H., Dahl, C. M., and Kristensen, J. T. (2011). Headlights on tobacco road to low birthweight outcomes — evidence from a battery of quantile regression estimators and a heterogeneous panel. CREATES Research Papers 2008-20, School of Economics and Management, University of Aarhus.

Barsky, R. B., Juster, F. T., Kimball, M. S., and Shapiro, M. D. (1997). Preference parameters and behavioral heterogeneity: An experimental approach in the Health and Retirement Study. *Quarterly Journal of Economics*, 112(2):537–79.

Bassett, Jr., G. and Koenker, R. (1982). An empirical quantile function for linear models with iid errors. *Journal of the American Statistical Association*, 77(378):407–15.

Beauchamp, J., Cesarini, D., and Johannesson, M. (2011). The psychometric properties of measures of economic risk preferences. Available on request via `https://files.nyu.edu/dac12/public/`.

Bellante, D. and Link, A. N. (1981). Are public sector workers more risk averse than private sector workers? *Industrial and Labor Relations Review*, 34(3):408–12.

Belzil, C. and Hansen, J. (2004). Earnings dispersion, risk aversion and education. In Polachek, S. W., editor, *Accounting for Worker Well-Being*, volume 23 of *Research in Labor Economics*, pages 335–58. Emerald Group Publishing Limited, Bingley, UK.

Belzil, C. and Leonardi, M. (2007). Can risk aversion explain schooling attainments? Evidence from Italy. *Labour Economics*, 14(6):957–70.

Belzil, C. and Leonardi, M. (2009). Risk aversion and schooling decisions. Working Papers hal-00411099, HAL.

Ben-Porath, Y. (1967). The production of human capital and the life cycle of earnings. *Journal of Political Economy*, 75(4):352–65.

Benjamin, D. J., Brown, S. A., and Shapiro, J. M. (2006). Who is "behavioral"? Cognitive ability and anomalous preferences. Available at `http://www.arts.cornell.edu/econ/dbenjamin/Who_is_Behavioral_Cognitive_ability_and_anomalous_preferences.pdf`.

Berkhout, P., Hartog, J., and Webbink, D. (2010). Compensation for earnings risk under worker heterogeneity. *Southern Economic Journal*, 76(3):762–90.

Bilias, Y., Chen, S., and Ying, Z. (2000). Simple resampling methods for censored regression quantiles. *Journal of Econometrics*, 99(2):373–86.

Bonin, H., Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2007). Cross-sectional earnings risk and occupational sorting: The role of risk attitudes. *Labour Economics*, 14(6):926–37.

Borghans, L., Golsteyn, B. H. H., Heckman, J. J., and Meijers, H. (2009). Gender differences in risk aversion and ambiguity aversion. *Journal of the European Economic Association*, 7(2-3):649–58.

Brown, C. and Moffitt, R. (1983). The effect of ignoring heteroscedasticity on estimates of the Tobit model. NBER Technical Working Papers 0027, National Bureau of Economic Research, Inc.

Brown, S., Farrell, L., Harris, M. N., and Sessions, J. G. (2006). Risk preference and employment contract type. *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 169(4):849–63.

Brunello, G. (2002). Absolute risk aversion and the returns to education. *Economics of Education Review*, 21(6):635–40.

Buchinsky, M. (1994). Changes in the U.S. wage structure 1963-1987: Application of quantile regression. *Econometrica*, 62(2):405–58.

Buchinsky, M. (1995). Estimating the asymptotic covariance matrix for quantile regression models: a Monte Carlo study. *Journal of Econometrics*, 68(2):303–38.

Buchinsky, M. (1998). Recent advances in quantile regression models: A practical guideline for empirical research. *Journal of Human Resources*, 33(1):88–126.

Buchinsky, M. and Hahn, J. (1998). An alternative estimator for the censored quantile regression model. *Econometrica*, 66(3):653–71.

Buckley, J. and James, I. (1979). Linear regression with censored data. *Biometrika*, 66(3):429–36.

Caliendo, M., Fossen, F., and Kritikos, A. (2009). Risk attitudes of nascent entrepreneurs: New evidence from an experimentally validated survey. *Small Business Economics*, 32(2):153–167.

Cameron, A. C. and Trivedi, P. K. (2005). *Microeconometrics*. Cambridge University Press, Cambridge, UK.

Cameron, A. C. and Trivedi, P. K. (2009). *Microeconometrics Using Stata*. StataCorp LP, College Station, TX.

Canay, I. A. (forthcoming). A simple approach to quantile regression for panel data. *Econometrics Journal*.

Caner, A. and Okten, C. (2010). Risk and career choice: Evidence from Turkey. *Economics of Education Review*, 29(6):1060–75.

Card, D. (1999). The causal effect of education on earnings. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3A, chapter 30, pages 1801–63. Elsevier, Amsterdam, The Netherlands.

Chamberlain, G. (1982). Multivariate regression models for panel data. *Journal of Econometrics*, 18(1):5–46.

Chamberlain, G. (1984). Panel data. In Griliches, Z. and Intriligator, M. D., editors, *Handbook of Econometrics*, volume 2, chapter 22, pages 1247–1318. Elsevier, Amsterdam, The Netherlands.

Chamberlain, G. (1994). Quantile regression, censoring, and the structuring of wages. In Sims, C. A., editor, *Advances in Econometrics: Sixth World Congress*, volume 1, chapter 5, pages 171–209. Cambridge University Press, Cambridge, UK.

Chay, K. Y. and Honoré, B. E. (1998). Estimation of semiparametric censored regression models: An application to changes in black-white earnings inequality during the 1960s. *Journal of Human Resources*, 33(1):4–38.

Chay, K. Y. and Powell, J. L. (2001). Semiparametric censored regression models. *Journal of Economic Perspectives*, 15(4):29–42.

Chen, S. H. (2008). Estimating the variance of wages in the presence of selection and unobserved heterogeneity. *The Review of Economics and Statistics*, 90(2):275–89.

Chernozhukov, V. (2005). Extremal quantile regression. *The Annals of Statistics*, 33(2):806–39.

Chernozhukov, V., Fernández-Val, I., and Galichon, A. (2010). Quantile and probability curves without crossing. *Econometrica*, 78(3):1093–1125.

Chernozhukov, V., Fernández-Val, I., Han, S., and Kowalski, A. E. (2011a). Stata command to implement CQIV. Available at `http://www.econ.yale.edu/~ak669/research.html`.

Chernozhukov, V., Fernández-Val, I., and Kowalski, A. E. (2011b). Quantile regression with censoring and endogeneity. Cowles Foundation Discussion Papers 1797, Cowles Foundation for Research in Economics, Yale University.

Chernozhukov, V. and Hong, H. (2002). Three-step censored quantile regression and extramarital affairs. *Journal of the American Statistical Association*, 97(459):872–82.

Chesher, A. and Irish, M. (1987). Residual analysis in the grouped and censored normal linear model. *Journal of Econometrics*, 34(1-2):33–61.

Cheung, G. (2007). Stability of the harm avoidance personality trait in late-life depression. *International Psychogeriatrics*, 19(4):778–80.

Cizek, P. (2008). Semiparametric robust estimation of truncated and censored regression models. Discussion Paper 2008-34, Tilburg University, Center for Economic Research.

Cloninger, C. R. (1987). A systematic method for clinical description and classification of personality variants. A proposal. *Archives of General Psychiatry*, 44(6):573–88.

Cox, J. C. and Oaxaca, R. L. (1989). Laboratory experiments with a finite-horizon job-search model. *Journal of Risk and Uncertainty*, 2(3):301–29.

Cramer, J. S., Hartog, J., Jonker, N., and van Praag, C. M. (2002). Low risk aversion encourages the choice for entrepreneurship: an empirical test of a truism. *Journal of Economic Behavior & Organization*, 48(1):29–36.

Cunha, F. and Heckman, J. J. (2007). Identifying and estimating the distributions of ex post and ex ante returns to schooling. *Labour Economics*, 14(6):870–93.

Dahl, C. M., le Maire, D., and Munch, J. R. (2009). Wage dispersion and decentralization of wage bargaining. Discussion Papers 09-15, University of Copenhagen. Department of Economics.

Di Mauro, C. and Musumeci, R. (forthcoming). Linking risk aversion and type of employment. *Journal of Socio-Economics*.

Diamond, J. (2003). *Guns, Germs, and Steel: The Fates of Human Societies*. W. W. Norton & Company, New York, NY, paperback edition.

Dohmen, T. and Falk, A. (2010). You get what you pay for: Incentives and selection in the education system. *Economic Journal*, 120(546):256–71.

Dohmen, T. and Falk, A. (2011). Performance pay and multidimensional sorting: Productivity, preferences, and gender. *American Economic Review*, 101(2):556–90.

Dohmen, T., Falk, A., Fliessbach, K., Sunde, U., and Weber, B. (2011a). Relative versus absolute income, joy of winning, and gender: Brain imaging evidence. *Journal of Public Economics*, 95(3-4):279–85.

Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review*, 100(3):1238–60.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., and Wagner, G. G. (2007). The measurement and stability of risk attitudes.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J., and Wagner, G. G. (2011b). Individual Risk Attitudes: Measurement, Determinants and Behavioral Consequences. *Journal of the European Economic Association*, 9(3):522–50.

Donald, S. G. (1995). Two-step estimation of heteroskedastic sample selection models. *Journal of Econometrics*, 65(2):347–80.

Drews, N. (2008). Das Regionalfile der IAB-Beschäftigtenstichprobe 1975-2004: Handbuch-Version 1.0.3. FDZ Datenreport 02/2008(DE), Institut für Arbeitsmarkt- und Berufsforschung (IAB), Nuremberg, Germany. Available at `http://doku.iab.de/fdz/reporte/2008/DR_02-08.pdf`.

Drukker, D. M. (2002). Bootstrapping a conditional moments test for normality after Tobit estimation. *Stata Journal*, 2(2):125–39.

Dustmann, C., Ludsteck, J., and Schönberg, U. (2009). Revisiting the German wage structure. *Quarterly Journal of Economics*, 124(2):843–81.

Eckel, C. C. and Grossman, P. J. (2008). *Men, Women and Risk Aversion: Experimental Evidence*, volume 1 of *Handbook of Experimental Economics Results*, chapter 113, pages 1061–73. Elsevier, Amsterdam, The Netherlands.

Ekelund, J., Johansson, E., Jarvelin, M.-R., and Lichtermann, D. (2005). Self-employment and risk aversion — evidence from psychological test data. *Labour Economics*, 12(5):649–59.

Fack, G. and Landais, C. (2010). Are tax incentives for charitable giving efficient? Evidence from France. *American Economic Journal: Economic Policy*, 2(2):117–41.

Federal Ministry of the Interior (2002). The public service in Germany. Available at `http://www.bmi.bund.de/SharedDocs/Downloads/EN/Broschueren/Der_oeffentliche_Dienst_in_Deutschland_Id_24276_en.pdf?__blob=publicationFile`.

Federal Ministry of the Interior (2009). The federal public service. Available at `http://www.bmi.bund.de/SharedDocs/Downloads/EN/Broschueren/2009/oed.pdf?__blob=publicationFile`.

Feinberg, R. M. (1977). Risk aversion, risk, and the duration of unemployment. *The Review of Economics and Statistics*, 59(3):264–71.

Feinberg, R. M. (1981). Earnings-risk as a compensating differential. *Southern Economic Journal*, 48(1):156–63.

Findeisen, S. (2010). Occupational choice, income risk and preference heterogeneity. Available at `http://www.iza.org/conference_files/SUMS2010/findeisen_s4081.pdf`.

Fitzenberger, B. (1994). A note on estimating censored quantile regressions. Technical Report Discussion Paper No. 14, Center of International Labor Economics, University of Konstanz.

Fitzenberger, B. (1997a). Computational aspects of censored quantile regression. *Lecture Notes-Monograph Series*, 31:171–86.

Fitzenberger, B. (1997b). A guide to censored quantile regressions. In Maddala, G. and Rao, C., editors, *Robust Inference*, volume 15 of *Handbook of Statistics*, pages 405–37. Elsevier, Amsterdam, The Netherlands.

Fitzenberger, B. (1998). The moving blocks bootstrap and robust inference for linear least squares and quantile regressions. *Journal of Econometrics*, 82(2):235–87.

Fitzenberger, B., Osikominu, A., and Völter, R. (2006). Imputation rules to improve the education variable in the IAB employment subsample. *Schmollers Jahrbuch*, 126(3):405–36.

Fouarge, D., Kriechel, B., and Dohmen, T. (2011). Occupational sorting of school leavers: The role of economic preferences.

Friedman, M. (1953). Choice, chance, and the personal distribution of income. *Journal of Political Economy*, 61(4):277–90.

Fuchs-Schündeln, N. and Schündeln, M. (2005). Precautionary savings and self-selection: Evidence from the German reunification experiment. *Quarterly Journal of Economics*, 120(3):1085–1120.

Gibbons, R. and Murphy, K. J. (1992). Optimal incentive contracts in the presence of career concerns: Theory and evidence. *Journal of Political Economy*, 100(3):468–505.

Goldberger, A. (1983). Abnormal selection bias. In Karlin, S., Amemiya, T., and Goodman, L., editors, *Studies in Econometrics, Time Series, and Multivariate Statistics*, pages 67–85. Academic Press, New York, NY.

Gourieroux, C., Monfort, A., Renault, E., and Trognon, A. (1987). Generalised residuals. *Journal of Econometrics*, 34(1-2):5–32.

Guiso, L., Jappelli, T., and Pistaferri, L. (2002). An empirical analysis of earnings and employment risk. *Journal of Business & Economic Statistics*, 20(2):241–53.

Guiso, L. and Paiella, M. (2005). The role of risk aversion in predicting individual behavior. Temi di discussione (Economic working papers) 546, Bank of Italy, Economic Research Department.

Gustavsen, G. W., Jolliffe, D., and Rickertsen, K. (2008). Censored quantile regression and purchases of ice cream. *Food Economics*, 5(3):152–63.

Harris, M. and Weiss, Y. (1984). Job matching with finite horizon and risk aversion. *Journal of Political Economy*, 92(4):758–79.

Hartog, J. (2009). A risk augmented Mincer earnings equation? Taking stock. IZA Discussion Papers 4439, Institute for the Study of Labor (IZA).

Hartog, J., Ferrer-i-Carbonell, A., and Jonker, N. (2002). Linking measured risk aversion to individual characteristics. *Kyklos*, 55(1):3–26.

Hartog, J., Plug, E., Diaz-Serrano, L., and Vieira, J. (2003). Risk compensation in wages — a replication. *Empirical Economics*, 28(3):639–47.

Hartog, J. and Vijverberg, W. P. M. (2007). On compensation for risk aversion and skewness affection in wages. *Labour Economics*, 14(6):938–56.

Heckman, J. J. (1976). The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models. In *Annals of Economic and Social Measurement*, volume 5 of *NBER Chapters*, pages 120–37. National Bureau of Economic Research, Inc.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–61.

Heckman, J. J., Lochner, L. J., and Todd, P. E. (2003). Fifty years of Mincer earnings regressions. NBER Working Papers 9732, National Bureau of Economic Research, Inc.

Heckman, J. J., Lochner, L. J., and Todd, P. E. (2006). Earnings functions, rates of return and treatment effects: The Mincer equation and beyond. volume 1 of *Handbook of the Economics of Education*, chapter 7, pages 307–458. Elsevier, Amsterdam, The Netherlands.

Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5):1644–55.

Honoré, B. E. (1992). Trimmed LAD and least squares estimation of truncated and censored regression models with fixed effects. *Econometrica*, 60(3):533–65.

Honoré, B. E. (1993). Orthogonality conditions for Tobit models with fixed effects and lagged dependent variables. *Journal of Econometrics*, 59(1-2):35–61.

Honoré, B. E. (2009). Pantob (version 0.6). Available at `http://www.princeton.edu/~honore/stata/index.html`.

Honoré, B. E. and Hu, L. (2004). Estimation of cross sectional and panel data censored regression models with endogeneity. *Journal of Econometrics*, 122(2):293–316.

Honoré, B. E., Khan, S., and Powell, J. L. (2002). Quantile regression under random censoring. *Journal of Econometrics*, 109(1):67–105.

Honoré, B. E. and Kyriazidou, E. (2000). Estimation of Tobit-type models with individual specific effects. *Econometric Reviews*, 19(3):341–66.

Hsiao, C. (2003). *Analysis of Panel Data*. Cambridge University Press, Cambridge, UK.

Hu, L. (2002). Estimation of a censored dynamic panel data model. *Econometrica*, 70(6):2499–2517.

Hurd, M. (1979). Estimation in truncated samples when there is heteroscedasticity. *Journal of Econometrics*, 11(2-3):247–58.

Isphording, I. E. (2010). Risky business — the role of individual risk attitudes in occupational choice. Ruhr Economic Papers 187, Ruhr-Universität Bochum (RUB), Department of Economics.

Jolliffe, D., Krushelnytskyy, B., and Semykina, A. (2000). Censored least absolute deviations estimator: CLAD. *Stata Technical Bulletin*, 58:13–16.

Kaplan, E. L. and Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association*, 53(282):457–81.

Khan, S. and Powell, J. L. (2001). Two-step estimation of semiparametric censored regression models. *Journal of Econometrics*, 103(1-2):73–110.

Kihlstrom, R. E. and Laffont, J.-J. (1979). A general equilibrium entrepreneurial theory of firm formation based on risk aversion. *Journal of Political Economy*, 87(4):719–48.

King, A. G. (1974). Occupational choice, risk aversion, and wealth. *Industrial and Labor Relations Review*, 27(4):586–96.

KMK (2010). The education system in the Federal Republic of Germany 2008 — a description of the responsibilities, structures and developments in education policy for the exchange of information in Europe. Available at `http://www.kmk.org/fileadmin/doc/Dokumentation/Bildungswesen_en_pdfs/dossier_en_ebook.pdf`.

Koenker, R. (2004). Quantile regression for longitudinal data. *Journal of Multivariate Analysis*, 91(1):74–89.

Koenker, R. (2005). *Quantile Regression*. Cambridge University Press, Cambridge, UK.

Koenker, R. and Bassett, Jr., G. (1978). Regression quantiles. *Econometrica*, 46(1):33–50.

Koenker, R. and Bassett, Jr., G. (1982). Robust tests for heteroscedasticity based on regression quantiles. *Econometrica*, 50(1):43–61.

Koenker, R. and Hallock, K. F. (2000). Quantile regression: an introduction. Technical report, University of Illinois at Urbana-Champaign.

Koenker, R. and Hallock, K. F. (2001). Quantile regression. *Journal of Economic Perspectives*, 15(4):143–56.

Koenker, R. and Xiao, Z. (2002a). Inference on the quantile regression process. *Econometrica*, 70(4):1583–1612.

Koenker, R. and Xiao, Z. (2002b). Inference on the quantile regression process: Electronic appendix.

Kowalski, A. E. (2009). Censored quantile instrumental variable estimates of the price elasticity of expenditure on medical care. NBER Working Papers 15085, National Bureau of Economic Research, Inc.

Lazear, E. P. and Shaw, K. L. (2009). Wage structure, raises and mobility: An introduction to international comparisons of the structure of wages within and across firms. In *The Structure of Wages: An International Comparison*, NBER Chapters. National Bureau of Economic Research, Inc.

Le, A. T., Miller, P. W., Slutske, W. S., and Martin, N. G. (2011a). Attitudes towards economic risk and occupational choice. Under review.

Le, A. T., Miller, P. W., Slutske, W. S., and Martin, N. G. (2011b). Attitudes towards economic risk and the gender pay gap. *Labour Economics*, 18(4):555–61.

Lehmer, F. and Möller, J. (2010). Interrelations between the urban wage premium and firm-size wage differentials: a microdata cohort analysis for Germany. *The Annals of Regional Science*, 45(1):31–53.

Leigh, J. P. (1981). Occupational choice under earnings uncertainty. *Nebraska Journal of Economics & Business*, 20:59–71.

Levhari, D. and Weiss, Y. (1974). The effect of risk on the investment in human capital. *American Economic Review*, 64(6):950–63.

Lönnqvist, J.-E., Verkasalo, M., Walkowitz, G., and Wichardt, P. C. (2011). Measuring individual risk attitudes in the lab: Task or ask? An empirical comparison. SOEPpapers 370, DIW Berlin, The German Socio-Economic Panel (SOEP).

Ludsteck, J. (2008). Wage cyclicality and the wage curve under the microscope. IAB Discussion Paper 200811, Institut für Arbeitsmarkts- und Berufsforschung (IAB), Nuremberg, Germany.

Ludsteck, J. and Haupt, H. (2007). An empirical test of the Reder hypothesis. Discussion Papers in Economics 1397, University of Munich, Department of Economics.

Machado, J. A. F. and Santos Silva, J. M. C. (2008). Quantiles for fractions and other mixed data. Economics Discussion Papers 656, University of Essex, Department of Economics.

Maddala, G. and Nelson, F. D. (1975). Specification errors in limited dependent variable models. NBER Working Papers 0096, National Bureau of Economic Research, Inc.

McCrae, R. R. (1993). Moderated analyses of longitudinal personality stability. *Journal of Personality and Social Psychology*, 65(3):577–85.

McGoldrick, K. (1995). Do women receive compensating wages for earnings uncertainty? *Southern Economic Journal*, 62(1):210–22.

McGue, M., Bacon, S., and Lykken, D. T. (1993). Personality stability and change in early adulthood: A behavioral genetic analysis. *Developmental Psychology*, 29(1):96–109.

Melenberg, B. and van Soest, A. (1996). Parametric and semi-parametric modelling of vacation expenditures. *Journal of Applied Econometrics*, 11(1):59–76.

Mincer, J. (1991). Education and unemployment. NBER Working Papers 3838, National Bureau of Economic Research, Inc.

Mincer, J. A. (1958). Investment in human capital and personal income distribution. *Journal of Political Economy*, 66(4):281–302.

Mincer, J. A. (1974). *Schooling, Experience, and Earnings*. Columbia University Press, New York, NY.

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica*, 46(1):69–85.

Newey, W. K. (1985). Maximum likelihood specification testing and conditional moment tests. *Econometrica*, 53(5):1047–70.

Newey, W. K. and Powell, J. L. (1990). Efficient estimation of linear and type I censored regression models under conditional quantile restrictions. *Econometric Theory*, 6(3):295–317.

Nickell, S. and Bell, B. (1996). Changes in the distribution of wages and unemployment in OECD countries. *The American Economic Review*, 86(2):302–8. Papers and Proceedings of the Hundredth and Eighth Annual Meeting of the American Economic Association, San Francisco, CA, January 5-7, 1996.

Paarsch, H. J. (1984). A Monte Carlo comparison of estimators for censored regression models. *Journal of Econometrics*, 24(1-2):197–213.

Pagan, A. and Vella, F. (1989). Diagnostic tests for models based on individual data: A survey. *Journal of Applied Econometrics*, 4(S):29–59.

Pfeifer, C. (2011). Risk aversion and sorting into public sector employment. *German Economic Review*, 12(1):85–99.

Pissarides, C. A. (1974). Risk, job search, and income distribution. *Journal of Political Economy*, 82(6):1255–67.

Polachek, S. W. (2003). Mincer's overtaking point and the lifecycle earnings distribution. IZA Discussion Papers 865, Institute for the Study of Labor (IZA).

Ponomareva, M. (2011). Quantile regression for panel data models with fixed effects and small T: Identification and estimation. Available at `http://publish.uwo.ca/~mponomar/QR.pdf`.

Powell, J. L. (1984). Least absolute deviations estimation for the censored regression model. *Journal of Econometrics*, 25(3):303–25.

Powell, J. L. (1986a). Censored regression quantiles. *Journal of Econometrics*, 32(1):143–55.

Powell, J. L. (1986b). Symmetrically trimmed least squares estimation for Tobit models. *Econometrica*, 54(6):1435–60.

Powell, J. L. (1994). Estimation of semiparametric models. In Engle, R. F. and McFadden, D., editors, *Handbook of Econometrics*, volume 4, chapter 41, pages 2443–2521. Elsevier, Amsterdam, The Netherlands.

Projektgruppe SOEP (2009). *Documentation PGEN — Person-related status and generated variables.* DIW, Berlin.

Rogers, W. (1993). Calculation of quantile regression standard errors. *Stata Technical Bulletin*, 13:18–19.

Rousseeuw, P. J. and Croux, C. (1993). Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424):1273–83.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers*, 3(2):135–46.

Ruud, P. (1984). Tests of specification in econometrics. *Econometric Reviews*, 3(2):211–42.

Sahm, C. R. (2008). How much does risk tolerance change? Finance and Economics Discussion Series 2007-66, Board of Governors of the Federal Reserve System (U.S.).

Saks, R. E. and Shore, S. H. (2005). Risk and career choice. *The B.E. Journal of Economic Analysis & Policy*, 5(1).

Schmillen, A. and Möller, J. (2010). Determinants of lifetime unemployment — a micro data analysis with censored quantile regressions. Available at `http://client.norc.org/jole/soleweb/11164.pdf`.

Schulhofer-Wohl, S. (2009). Heterogeneity and tests of risk sharing. Available at `http://econ.econ.duke.edu/uploads/assets/Workshop%20Papers/schulhofer-wohl.pdf`.

Shaw, K. L. (1996). An empirical analysis of risk aversion and income growth. *Journal of Labor Economics*, 14(4):626–53.

Skeels, C. L. and Vella, F. (1999). A Monte Carlo investigation of the sampling behavior of conditional moment tests in Tobit and Probit models. *Journal of Econometrics*, 92(2):275–94.

Smith, V. K. (1973). Least squares regression with Cauchy errors. *Oxford Bulletin of Economics and Statistics*, 35(3):223–31.

Statistisches Bundesamt Deutschland (2011). Verbraucherpreise. Available at `http://www.destatis.de/jetspeed/portal/cms/Sites/destatis/Internet/EN/Content/Statistics/TimeSeries/EconomicIndicators/BasicData/Content100/vpi101a,templateId=renderPrint.psml`.

Tobin, J. (1958). Estimation of relationships for limited dependent variables. *Econometrica*, 26(1):24–36.

Vijverberg, W. P. M. (1987). Non-normality as distributional misspecification in single-equation limited dependent variable models. *Oxford Bulletin of Economics and Statistics*, 49(4):417–30.

Wagner, G. G., Frick, J. R., and Schupp, J. (2007). The German Socio-Economic Panel Study (SOEP) — Scope, Evolution and Enhancements. *Schmollers Jahrbuch*, 127(1):139–69.

Weisbrod, B. A. (1962). Education and investment in human capital. *Journal of Political Economy*, 70(5):106–23.

Weiss, A. A. (1991). Estimating nonlinear dynamic models using least absolute error estimation. *Econometric Theory*, 7(1):46–68.

Wilhelm, M. O. (2008). Practical considerations for choosing between Tobit and SCLS or CLAD estimators for censored regression models with an application to charitable giving. *Oxford Bulletin of Economics and Statistics*, 70(4):559–82.

A Statistics and results

A.1 Degree of censoring

Table 1: Degree of censoring by groups of age and education

	25-34	35-44	45-54	Total
No degree	0.4%	1.2%	1.0%	0.9%
Vocational training (VT)	0.9%	3.8%	6.2%	3.9%
High school (HS)	7.2%	18.4%	15.2%	12.4%
HS & VT	4.9%	16.9%	20.0%	13.5%
Technical college	11.0%	35.1%	43.4%	32.1%
University	20.0%	47.5%	55.5%	43.8%
Total	3.6%	12.5%	13.7%	10.5%

Fractions of censored wages in the IAB Employment Sample. West German men in full-time employment — June 30, 2004. Education corrected as discussed in section 3.2.

A.2 Mincer estimates

Please turn page.

Table 2: Human capital model

Dependent variable: log of censored daily wage						
	(1)	(2)	(3)	(4)	(5)	(6)
Experience	0.064***	0.058***	0.060***	0.058***	0.057***	0.050***
	(0.002)	(0.006)	(0.005)	(0.003)	(0.004)	(0.005)
Experience^2/100	-0.236***	-0.228***	-0.224***	-0.211***	-0.202***	-0.170***
	(0.008)	(0.031)	(0.024)	(0.016)	(0.020)	(0.027)
Experience^3/1000	0.029***	0.029***	0.028***	0.026***	0.024***	0.020***
	(0.001)	(0.005)	(0.004)	(0.003)	(0.003)	(0.004)
Tenure	0.025***	0.039***	0.027***	0.020***	0.018***	0.016***
	(0.000)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Tenure^2/100	-0.057***	-0.091***	-0.061***	-0.043***	-0.038***	-0.035***
	(0.001)	(0.005)	(0.003)	(0.003)	(0.003)	(0.005)
Years of Education	0.056***	0.049***	0.050***	0.052***	0.052***	0.049***
	(0.000)	(0.002)	(0.002)	(0.001)	(0.002)	(0.003)
Constant	3.305***	3.004***	3.203***	3.419***	3.666***	3.952***
	(0.012)	(0.050)	(0.037)	(0.026)	(0.036)	(0.068)
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	168,863	168,863	168,863	168,863	168,863	168,863

Human capital model estimated using (1) Tobit ML approach, and (2)-(6) the three-step CQR estimator at the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th} percentile. Standard errors of coefficient estimates, computed using direct percentile (bootstrap) method and asymptotic normality, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable is the log of censored daily wage. Occupation dummies based on IABS KldB 88 coding.

A.3 Replication of Bonin et al. (2007)

Please turn page.

Table 3: Replication of Bonin et al. (2007)

Dependent variable: 10-90 spread of occupation effects						
	(1)	(2)	(3)	(4)	(5)	(6)
General Risk Attitude	0.002***	0.002***	0.001**			
	(0.001)	(0.001)	(0.001)			
General Risk Attitude (av.)				0.003***	0.002**	0.002**
				(0.001)	(0.001)	(0.001)
Experience		0.001***	0.000		0.001***	0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Tenure		0.000	-0.000		0.000	-0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Years of Education		0.008**	0.006*		0.008**	0.006*
		(0.004)	(0.003)		(0.004)	(0.003)
Married and living together		-0.006	-0.011*		-0.006	-0.011*
		(0.005)	(0.006)		(0.005)	(0.006)
Body height		0.001***	0.001***		0.001***	0.001***
		(0.000)	(0.000)		(0.000)	(0.000)
East Germany		-0.013*	-0.001		-0.013*	-0.001
		(0.006)	(0.007)		(0.006)	(0.007)
Public Sector Employment		0.027*	0.032**		0.027**	0.033**
		(0.013)	(0.015)		(0.013)	(0.015)
Log Monthly Income			0.036**			0.036**
			(0.016)			(0.016)
Constant	-0.067***	-0.332***	-0.558***	-0.068***	-0.336***	-0.561***
	(0.014)	(0.075)	(0.157)	(0.014)	(0.075)	(0.156)
Observations	4,042	3,941	3,566	4,052	3,951	3,572
R-squared	0.004	0.114	0.149	0.003	0.115	0.149

OLS estimates. Robust standard errors of coefficient estimates, allowing for clustering at the ISCO 2-digit occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable is the 10-90 spread of occupation dummy estimates from QR, on ISCO 2-digit level. "General Risk Attitude" is the response to the 2004 general risk question, "General Risk Attitude (av.)" is the corresponding average over the 2004, 2006, 2008, and 2009 responses.

Table 4: Replication of Bonin et al. (2007)

Dependent variable: median differences of occupation effects								
	10-50		25-50		50-75		50-90	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
General Risk Attitude	0.001**		0.001**		0.000		0.001	
	(0.000)		(0.000)		(0.000)		(0.001)	
General Risk Attitude (av.)		0.000		0.001**		0.000		0.001*
		(0.000)		(0.000)		(0.000)		(0.001)
Experience	0.000	0.000	-0.000	-0.000	-0.000	-0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Tenure	-0.000	-0.000	-0.000	-0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Years of Education	0.002	0.002	0.003**	0.003**	0.002	0.002	0.004	0.004
	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)	(0.003)	(0.003)
Married and living together	-0.003	-0.003	-0.003*	-0.003*	-0.003*	-0.003*	-0.007	-0.007
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.005)	(0.005)
Body height	0.000	0.000	0.000	0.000	0.000**	0.000**	0.000**	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
East Germany	0.005	0.005	0.002	0.002	-0.002	-0.002	-0.006	-0.006
	(0.005)	(0.005)	(0.003)	(0.003)	(0.002)	(0.002)	(0.004)	(0.004)
Public Sector Employment	0.004	0.004	0.009	0.010	0.008	0.008	0.029**	0.029**
	(0.010)	(0.010)	(0.008)	(0.008)	(0.006)	(0.006)	(0.013)	(0.013)
Log Monthly Income	0.027**	0.027**	0.025***	0.025***	0.016***	0.016***	0.009	0.009
	(0.011)	(0.011)	(0.008)	(0.008)	(0.005)	(0.005)	(0.009)	(0.009)
Constant	-0.225*	-0.227*	-0.258***	-0.259***	-0.238***	-0.238***	-0.332***	-0.334***
	(0.113)	(0.114)	(0.076)	(0.076)	(0.052)	(0.052)	(0.103)	(0.103)
Observations	3,566	3,572	3,566	3,572	3,566	3,572	3,566	3,572
R-squared	0.090	0.090	0.196	0.196	0.168	0.169	0.113	0.114

OLS estimates. Robust standard errors of coefficient estimates, allowing for clustering at the ISCO 2-digit occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable are the median differences of occupation dummy estimates from QR, on ISCO 2-digit level. "General Risk Attitude" is the response to the 2004 general risk question, "General Risk Attitude (av.)" is the corresponding average over the 2004, 2006, 2008, and 2009 responses.

A.4 IABS risk profiles

Please turn page.

Table 5: IABS risk profiles

Dependent variable: 10-90 spread of occupation effects						
	(1)	(2)	(3)	(4)	(5)	(6)
General Risk Attitude	0.003**	0.002*	0.003*			
	(0.002)	(0.001)	(0.002)			
General Risk Attitude (av.)				0.005**	0.005**	0.005**
				(0.002)	(0.002)	(0.002)
Experience		0.000	0.000		0.000	0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Tenure		-0.000	0.000		-0.000	0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Years of Education		0.002	0.010**		0.002	0.010**
		(0.004)	(0.004)		(0.004)	(0.004)
Married and living together		-0.017**	-0.016**		-0.017**	-0.017**
		(0.007)	(0.007)		(0.007)	(0.007)
Body height		0.000	0.001		0.000	0.001
		(0.000)	(0.000)		(0.000)	(0.000)
Public Sector Employment		-0.015	-0.017		-0.014	-0.016
		(0.026)	(0.026)		(0.026)	(0.026)
Median wage (occ.)			-0.122*			-0.122*
			(0.068)			(0.068)
Constant	-0.188***	-0.206**	0.164	-0.198***	-0.221**	0.150
	(0.014)	(0.087)	(0.264)	(0.014)	(0.086)	(0.261)
Observations	2,815	2,740	2,740	2,822	2,747	2,747
R-squared	0.002	0.009	0.051	0.004	0.011	0.054

OLS estimates. Robust standard errors of coefficient estimates, allowing for clustering at the IABS KldB 88 occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable is the 10-90 spread of occupation dummy estimates from QR, on IABS KldB 88 level. "General Risk Attitude" is the response to the 2004 general risk question, "General Risk Attitude (av.)" is the corresponding average over the 2004, 2006, 2008, and 2009 responses.

Table 6: IABS risk profiles

Dependent variable: 25-75 spread of occupation effects						
	(1)	(2)	(3)	(4)	(5)	(6)
General Risk Attitude	0.002**	0.001	0.001			
	(0.001)	(0.001)	(0.001)			
General Risk Attitude (av.)				0.003***	0.003**	0.003***
				(0.001)	(0.001)	(0.001)
Experience		-0.000	-0.000		-0.000	-0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Tenure		0.000	0.000		0.000	0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Years of Education		0.004	0.005*		0.004	0.005*
		(0.003)	(0.003)		(0.003)	(0.003)
Married and living together		-0.009*	-0.007		-0.009*	-0.008
		(0.005)	(0.005)		(0.005)	(0.005)
Body height		0.000	0.000		0.000	0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Public Sector Employment		-0.010	-0.010		-0.010	-0.009
		(0.015)	(0.015)		(0.015)	(0.015)
Median wage (occ.)			-0.012			-0.012
			(0.041)			(0.041)
Constant	-0.097***	-0.187***	-0.163	-0.104***	-0.192***	-0.171
	(0.009)	(0.061)	(0.169)	(0.008)	(0.061)	(0.167)
Observations	2,815	2,740	2,740	2,822	2,747	2,747
R-squared	0.003	0.025	0.026	0.005	0.027	0.029

OLS estimates. Robust standard errors of coefficient estimates, allowing for clustering at the IABS KldB 88 occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable is the 25-75 spread of occupation dummy estimates from QR, on IABS KldB 88 level. "General Risk Attitude" is the response to the 2004 general risk question, "General Risk Attitude (av.)" is the corresponding average over the 2004, 2006, 2008, and 2009 responses.

Table 7: IABS risk profiles

Dependent variable: median differences of occupation effects								
	10-50		25-50		50-75		50-90	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
General Risk Attitude	0.002*		0.000		0.001**		0.002*	
	(0.001)		(0.000)		(0.000)		(0.001)	
General Risk Attitude (av.)		0.003**		0.001*		0.002***		0.003**
		(0.001)		(0.001)		(0.001)		(0.001)
Experience	0.000	0.000	0.000	0.000	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Tenure	-0.000	-0.000	0.000	0.000	0.000**	0.000**	0.000*	0.000*
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Years of Education	0.005*	0.005*	0.002*	0.002*	0.003*	0.003*	0.005**	0.005**
	(0.003)	(0.003)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Married and living together	-0.009**	-0.009**	-0.003	-0.004	-0.005*	-0.005*	-0.008**	-0.008**
	(0.004)	(0.004)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)
Body height	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Public Sector Employment	-0.000	-0.000	-0.009	-0.009	-0.001	-0.001	-0.016	-0.016
	(0.016)	(0.016)	(0.007)	(0.007)	(0.009)	(0.009)	(0.012)	(0.012)
Median wage (occ.)	-0.012	-0.012	-0.009	-0.009	-0.002	-0.002	-0.110***	-0.110***
	(0.048)	(0.048)	(0.022)	(0.022)	(0.022)	(0.022)	(0.032)	(0.032)
Constant	-0.114	-0.123	-0.046	-0.049	-0.109	-0.112	0.278**	0.273**
	(0.209)	(0.208)	(0.096)	(0.095)	(0.084)	(0.083)	(0.122)	(0.121)
Observations	2,740	2,747	2,740	2,747	2,740	2,747	2,740	2,747
R-squared	0.025	0.027	0.020	0.022	0.032	0.034	0.133	0.135

OLS estimates. Robust standard errors of coefficient estimates, allowing for clustering at the IABS KldB 88 occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable are the median differences of occupation dummy estimates from QR, on IABS KldB 88 level. "General Risk Attitude" is the response to the 2004 general risk question, "General Risk Attitude (av.)" is the corresponding average over the 2004, 2006, 2008, and 2009 responses.

Table 8: IABS risk profiles

Dependent variable: 10-90 spread of occupation effects						
	(1)	(2)	(3)	(4)	(5)	(6)
Career Risk Attitude	0.004**	0.003**	0.004**			
	(0.002)	(0.001)	(0.002)			
Career Risk Attitude (av.)				0.004**	0.003**	0.004**
				(0.002)	(0.001)	(0.002)
Experience		0.000	0.000		0.000	0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Tenure		-0.000	0.000		-0.000	0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Years of Education		0.002	0.010**		0.002	0.010**
		(0.004)	(0.004)		(0.004)	(0.004)
Married and living together		-0.017**	-0.016**		-0.017**	-0.017**
		(0.007)	(0.007)		(0.007)	(0.007)
Body height		0.000	0.001		0.000	0.001
		(0.000)	(0.000)		(0.000)	(0.000)
Public Sector Employment		-0.015	-0.017		-0.014	-0.016
		(0.026)	(0.026)		(0.026)	(0.026)
Median wage (occ.)			-0.125*			-0.125*
			(0.069)			(0.069)
Constant	-0.188***	-0.210**	0.169	-0.187***	-0.212**	0.167
	(0.014)	(0.087)	(0.265)	(0.014)	(0.087)	(0.265)
Observations	2,806	2,731	2,731	2,816	2,741	2,741
R-squared	0.004	0.010	0.054	0.003	0.010	0.054

OLS estimates. Robust standard errors of coefficient estimates, allowing for clustering at the IABS KldB 88 occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable is the 10-90 spread of occupation dummy estimates from QR, on IABS KldB 88 level. "Career Risk Attitude" is the response to the 2004 career risk question, "Career Risk Attitude (av.)" is the corresponding average over the 2004 and 2009 responses.

Table 9: IABS risk profiles

Dependent variable: median differences of occupation effects								
	10-50		25-50		50-75		50-90	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Career Risk Attitude	0.002**		0.001**		0.001**		0.002**	
	(0.001)		(0.000)		(0.000)		(0.001)	
Career Risk Attitude (av.)		0.002**		0.001**		0.001***		0.002**
		(0.001)		(0.000)		(0.001)		(0.001)
Experience	0.000	0.000	0.000	0.000	-0.000	-0.000	-0.000	-0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Tenure	-0.000	-0.000	0.000	0.000	0.000**	0.000**	0.000**	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Years of Education	0.005*	0.005*	0.002*	0.002*	0.003*	0.003*	0.005**	0.005**
	(0.003)	(0.003)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)
Married and living together	-0.009**	-0.009**	-0.003	-0.003	-0.005*	-0.005*	-0.008**	-0.008**
	(0.004)	(0.004)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)
Body height	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Public Sector Employment	-0.000	-0.000	-0.009	-0.009	-0.001	-0.001	-0.016	-0.016
	(0.016)	(0.016)	(0.007)	(0.007)	(0.009)	(0.010)	(0.012)	(0.012)
Median wage (occ.)	-0.013	-0.014	-0.009	-0.009	-0.002	-0.002	-0.112***	-0.111***
	(0.048)	(0.048)	(0.022)	(0.022)	(0.022)	(0.022)	(0.032)	(0.032)
Constant	-0.112	-0.114	-0.047	-0.047	-0.108	-0.107	0.281**	0.282**
	(0.209)	(0.208)	(0.094)	(0.094)	(0.084)	(0.085)	(0.123)	(0.124)
Observations	2,731	2,741	2,731	2,741	2,731	2,741	2,731	2,741
R-squared	0.027	0.028	0.023	0.023	0.035	0.034	0.135	0.134

OLS estimates. Robust standard errors of coefficient estimates, allowing for clustering at the IABS KldB 88 occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable are the median differences of occupation dummy estimates from QR, on IABS KldB 88 level. "Career Risk Attitude" is the response to the 2004 career risk question, "Career Risk Attitude (av.)" is the corresponding average over the 2004 and 2009 responses.

Table 10: IABS risk profiles — original education variable

Dependent variable: 25-75 spread of occupation effects						
	(1)	(2)	(3)	(4)	(5)	(6)
General Risk Attitude	0.002**	0.002*	0.001			
	(0.001)	(0.001)	(0.001)			
General Risk Attitude (av.)				0.004***	0.003**	0.002**
				(0.001)	(0.001)	(0.001)
Experience		-0.000	-0.000		-0.000	-0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Tenure		0.001*	0.001*		0.001*	0.001*
		(0.000)	(0.000)		(0.000)	(0.000)
Years of Education		0.007***	0.005*		0.007***	0.005*
		(0.002)	(0.002)		(0.002)	(0.002)
Married and living together		-0.008	-0.009*		-0.009*	-0.009*
		(0.005)	(0.005)		(0.005)	(0.005)
Body height		0.000	0.000		0.000	0.000
		(0.000)	(0.000)		(0.000)	(0.000)
Public Sector Employment		-0.003	-0.003		-0.003	-0.002
		(0.021)	(0.020)		(0.021)	(0.020)
Median wage (occ.)			0.041			0.042
			(0.038)			(0.038)
Constant	-0.095***	-0.244***	-0.369**	-0.101***	-0.248***	-0.374**
	(0.009)	(0.054)	(0.146)	(0.009)	(0.054)	(0.145)
Observations	2,815	2,740	2,740	2,822	2,747	2,747
R-squared	0.004	0.065	0.079	0.005	0.067	0.081

OLS estimates. Same as table 6, but using the original uncorrected education variable in the first-stage Mincer regressions. Robust standard errors of coefficient estimates, allowing for clustering at the IABS KldB 88 occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable is the 25-75 spread of occupation dummy estimates from QR, on IABS KldB 88 level. "General Risk Attitude (av.)" is the average over the 2004, 2006, 2008, and 2009 responses to the general risk question.

A.5 Panel risk

Please turn page.

Table 11: Panel risk

Dependent variable: 10-90 spread of occupation effects						
	(1)	(2)	(3)	(4)	(5)	(6)
	(1)	(2)	(3)	(4)	(5)	(6)
General Risk Attitude (av.)	-0.001**	-0.001**	-0.000	-0.000	-0.001***	-0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Experience		-0.000		0.000		-0.000**
		(0.000)		(0.000)		(0.000)
Tenure		-0.000***		-0.000***		-0.000***
		(0.000)		(0.000)		(0.000)
Years of Education		-0.002***		-0.001***		-0.002***
		(0.000)		(0.000)		(0.000)
Married and living together		0.004***		0.005**		0.004***
		(0.001)		(0.002)		(0.001)
Body height		-0.000*		-0.000		-0.000***
		(0.000)		(0.000)		(0.000)
East Germany						
Public Sector Employment		0.002		0.006***		-0.004***
		(0.001)		(0.002)		(0.001)
Constant	0.034***	0.081***	0.034***	0.056**	0.034***	0.111***
	(0.002)	(0.014)	(0.003)	(0.022)	(0.002)	(0.015)
Observations	3,622	3,531	1,802	1,759	1,820	1,772
R-squared	0.002	0.029	0.000	0.018	0.010	0.101

OLS estimates. Robust standard errors of coefficient estimates, allowing for clustering at the ISCO 2-digit occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable is the 10-90 spread of occupation dummy estimates from QR after elimination of fixed effects, on ISCO 2-digit level. "General Risk Attitude (av.)" is the average over the 2004, 2006, 2008, and 2009 responses to the general risk question. (1), (2) all individuals; (3), (4) individuals with no prior occupation change; (5), (6) individuals with prior occupation change.

Table 12: Panel risk — switchers only

Dependent variable: median differences of occupation effects								
	10-50		25-50		50-75		50-90	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
General Risk Attitude (av.)	-0.001***	-0.001***	-0.001***	-0.000***	-0.000***	-0.000***	-0.001***	-0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Experience		-0.000		-0.000		-0.000**		-0.000***
		(0.000)		(0.000)		(0.000)		(0.000)
Tenure		-0.000***		-0.000***		-0.000		-0.000
		(0.000)		(0.000)		(0.000)		(0.000)
Years of Education		-0.002***		-0.001***		-0.001***		-0.001***
		(0.000)		(0.000)		(0.000)		(0.000)
Married and living together		0.002**		0.001**		0.001***		0.002***
		(0.001)		(0.000)		(0.000)		(0.001)
Body height		-0.000**		-0.000***		-0.000***		-0.000***
		(0.000)		(0.000)		(0.000)		(0.000)
East Germany								
Public Sector Employment		-0.004***		-0.003***		-0.000		0.000
		(0.001)		(0.000)		(0.000)		(0.001)
Constant	0.015***	0.060***	0.019***	0.050***	0.006***	0.033***	0.019***	0.051***
	(0.001)	(0.010)	(0.001)	(0.005)	(0.001)	(0.005)	(0.001)	(0.007)
Observations	1,820	1,772	1,820	1,772	1,820	1,772	1,820	1,772
R-squared	0.007	0.122	0.011	0.183	0.011	0.118	0.009	0.056

OLS estimates. Robust standard errors of coefficient estimates, allowing for clustering at the ISCO 2-digit occupation level, in parentheses; ***/**/* indicate significance at 1%/5%/10% level. Dependent variable are the median differences of occupation dummy estimates from QR after elimination of fixed effects, on ISCO 2-digit level. "General Risk Attitude" is the response to the 2004 general risk question, "General Risk Attitude (av.)" is the corresponding average over the 2004, 2006, 2008, and 2009 responses.

A.6 Stationarity of risk profiles

Table 13: Temporal stability of risk profiles

	1979	1984	1989	1994	1999	2004
1979	1.000					
1984	0.818	1.000				
1989	0.768	0.902	1.000			
1994	0.781	0.854	0.874	1.000		
1999	0.654	0.778	0.810	0.899	1.000	
2004	0.655	0.694	0.753	0.896	0.922	1.000

Pearson correlation coefficients of 10-90 spread per occupation, estimated as in section 6.2, across years.

B Monte Carlo evidence on estimation of dispersion

In this section, we review the performance of the estimation method described in section 5.2 and compare it to a residual-based method. For the former, we use the difference between the coefficient estimates of the group dummies at the 90^{th} and 10^{th} percentile; in the latter case, we estimate a conventional OLS regression including group dummies and compute the standard deviation of residuals per group. After each of 1,000 simulations, we compute the correlation of an occupation-specific scale σ_k and the two statistics. The models investigated are stylized versions of the wage distribution setting in our empirical analysis; specifically, we first consider a model with only group-specific scale, and then turn to location-scale models in which a regressor has a heteroskedastic effect.

B.1 Group-specific scale model

The DGP has the following linear representation:

$$Y_{ik} = \beta X_i + c_k + \sigma_k \epsilon_i \tag{39}$$

for $i \in \{1, \ldots, N\}$ and $k \in \{1, \ldots, m\}$, $c_k \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma^2)$, $\epsilon_i \overset{\text{iid}}{\sim} \mathcal{N}(0, 1)$.

In our simulation, we set $N = 1,000$, $X_i \overset{\text{iid}}{\sim} \mathcal{U}(0, 1)$, $\beta = .5$, $\sigma = .2$, and $\sigma_k = .1 + .2U_k$, where $U_k \overset{\text{iid}}{\sim} \mathcal{U}(0, 1)$. Individuals are randomly assigned to one of $m = 10$ groups according to a uniform distribution.

Table 14 shows a very similar performance for both statistics, with a correlation close to unity in each case.

Table 14: Group-specific scale model

	Mean	Std. Dev.	Min.	Max.
10-90 spread	0.974	0.023	0.738	0.999
Resid. std. dev.	0.983	0.015	0.861	0.999
Correlation of risk measures with true dispersion				

Table 15: Linear location-scale model

	Mean	Std. Dev.	Min.	Max.
10-90 spread	0.900	0.068	0.263	0.993
Resid. std. dev.	0.856	0.104	-0.083	0.992
Correlation of risk measures with true dispersion				

B.2 Linear location-scale model

Leaving all else the same,

$$Y_{ik} = \beta X_i + c_k + (\gamma X_i + \sigma_k)\epsilon_i \tag{40}$$

with $\gamma = .5$. Hence, the independent variable X now exerts a heteroskedastic effect.

Already, the 10-90 spread does slightly better (table 15).

B.3 Nonlinear location-scale model

We adapt the DGP such that the scale effect of the independent variable X is now negatively related to the occupation variance:

$$Y_{ik} = \beta X_i + c_k + [(1 - \delta\sigma_k)X_i + \sigma_k]\epsilon_i \tag{41}$$

For $\delta = 1$,

$$Y_{ik} = \beta X_i + c_k + [X_i + (1 - X_i)\sigma_k]\epsilon_i \tag{42}$$

As reported in table 16, the statistic based on QR is a lot more robust in this case, since the scale effect of X at the different quantiles is explicitly controlled for. The discrepancy will likely be even larger for more irregular distributions.

Table 16: Nonlinear location-scale model

	Mean	Std. Dev.	Min.	Max.
10-90 spread	0.824	0.115	0.295	0.991
Resid. std. dev.	0.676	0.208	-0.284	0.977
Correlation of risk measures with true dispersion				

C Imputation of years of schooling

C.1 SOEP

Like Bonin et al. (2007), we use the variable for years of education included in the SOEP (*$BILZEIT*); see Projektgruppe SOEP (2009).

C.2 IABS

Based on the categorical education variable *bild* of the IABS, we attempt to match the SOEP coding as closely as possible:

1. no degree at all: 10 years
2. vocational training degree: 11.5 years
3. high school degree: 13 years
4. high school and vocational training degree: 14.5 years
5. technical college degree: 16 years
6. university degree: 18 years