Is Self–Employment Really a Bad Experience?
The Effects of Previous Self–Employment on Subsequent Wage–Employment

Abstract: We use propensity score matching methods to quantify the effects of past self–employment experience on subsequent earnings in dependent employment using data on the population of Danish men observed between 1990 and 1996. Our results generally confirm existing studies in that we find that a spell of self–employment is associated with lower hourly wages compared to workers who were consecutively wage–employed. We also show, however, that this effect disappears — and even becomes positive in some settings — for formerly self–employed who find dependent employment in the same sector as their self–employment sector. Hence, the on average negative effect of self–employment is rather caused by sector–switching than by the self–employment experience *per se*. Moreover, formerly self–employed who either enjoyed a high income or hired at least one worker during their self–employment spell receive wages in subsequent dependent employment that are at least as high as for individuals who have been consecutively wage–employed.
JEL classification: $L26, C21$
Keywords: self–employment, entrepreneurship, wages, experience, propensity score matching

1 Introduction

While the benefits of self–employment for society are widely recognized (Blanchflower, 2000; OECD 2000, Ch. 5) and the value of "being your own boss" is undisputed (Benz and Frey, 2004; Hamilton, 2000), it is less clear what the monetary value of self–employment experience is for the wide majority of entrepreneurs who return to wage work after a spell of self–employment. From a theoretical point of view, experience from self–employment may provide either more or less (relevant) human capital than experience from wage employment.

This paper aims at finding empirical answers to the following questions: (1.) What is the average (or likely) effect of a previous self–employment spell on the subsequent wage rate in later wage–employment?, (2.) If there indeed is a wage difference, does it matter if the former self–employed changes the sector of employment?, and (3.) Does hiring workers or enjoying a high income during the self–employment spell matter for subsequent wage–employment wages?

The few empirical studies in this area either find negative or statistically insignificant effects of past self–employment on current wages in dependent employment. As probably the first contribution, Evans and Leighton (1989) do not find clear evidence of a different return to self–employment compared to wage–work experience for the US. Ferber and Waldfogel (1998) use a US sample of both self–employed and wage employed individuals and find a negative effect of self–employment experience from an unincorporated business when controlling for current self–employment, which in itself has a positive effect on returns. Using the same data set, Williams (2000) finds that the effect of self–employment experience on current wages is smaller than the effect of wage–work experience, but only significantly so for women. In a later study, Williams (2002) estimates the wage return to previous wage work and self–employment experience on German data and finds that self–employment experience yields a lower return in wage–employment than continued wage–employment. This result is shared by a recent analysis for Germany by Niefert (2006). Finally, Bruce and Schuetze (2004) find that brief spells of self–employment do not increase, and in some of their specifications significantly decrease, subsequent wage earnings compared to continued wage work for a sample of US workers.

The existing empirical evidence does tend to suffer, however, from imprecise estimates due to relatively small sample sizes and from econometric techniques that might not be able to correctly identify the effect of interest as discussed in more detail below. Additionally, none of these studies takes into account the differential effect of sector switching after a completed self–employment spell, or the characteristics of the self–employment spell as measured by the number of employees or the income while self–employed. However, according to both human–capital theory and signalling theory, both the sectoral transition pattern and the type of self–employment experience are likely to matter for the resulting effects of self–employment spells on subsequent wage–employment wages.

To analyze the effects of self–employment experience on subsequent wages in dependent employment, we use register data on the entire population of Danish male citizens observed between 1990 and 1996. The data set is particularly rich both in terms of information on worker characteristics and the number of workforce members it covers. The richness of our data allows us to perform a much more detailed and disaggregate analysis as well as to produce much more precise parameter estimates than existing studies.

We focus on male workers who have been wage employed both in 1990 and 1996 and analyze how a completed spell of self–employment within the five years 1991–1995 affects wage–employment wages in 1996. Our interest is hence *not* in *general* effects of self–employment but rather on the effects of previous self–employment spells on wages in subsequent wage–employment since, as we document in Section 3, self–employment spells tend to be short.

The reason for focusing on workers who have been wage–employed in both 1990 and 1996 is that we want to generate a homogeneous set of observations (all individuals are wage–employed both at the beginning and the end of the observation period) that we track for a relevant time period which is neither too short for self–employment effects to occur nor too long for self–employment effects to have faded out already. Considering a five–years' window is also consistent with Bruce and Schuetze (2004), a main reference of our paper.

We estimate the "average treatment effect on the treated" (ATT), where the treatment is the self–employment spell and the outcome variable is the subsequent wage in wage–employment in 1996. The ATT is not trivial to estimate in the present context since self–employed individuals may have observed and unobserved characteristics that are inherently different from those who are permanently wage–employed and since these characteristics may simultaneously affect the self–employment choice and wages. For example, if only high ability individuals become self–employed, assessing the wage differential without considering such an important characteristic would bias our results upward since high ability then influences both the selection into self–employment and wages positively. Not accounting for these differences between treatment and control group individuals then leads to a selection bias.

We use propensity score matching methods (PSM) to deal with this self–

selection problem. Our approach is to compare each individual that was self–employed between 1990 and 1996 to an individual who was consecutively wage–employed and who has the same observed characteristics. That way, we make sure that treatment and control group observations do not differ from one another with respect to core variables that are likely to affect both selection and 1996 wages, e.g. schooling, tenure, age, sector of employment, initial wage etc. Clearly, PSM does not account for unobserved characteristics that may differ between treatment and control group observations. This means that we properly identify causal self–employment effects conditional on observable characteristics only. At the same time, however, unobserved heterogeneity is likely to be reflected in observed characteristics that we do control for. For example, more able individuals are more likely to have enjoyed a longer education and a higher initial wage than less able ones and individuals who appreciate a stable working environment will have acquired more years of tenure than people with a taste for change.

We find that self–employment experience is, in general, valued less than continued wage work experience by subsequent employers. The average length of a self-employment spell in our sample is 1.8 years, and it is on average associated with a subsequent 2.9% loss in hourly wages compared to a situation where the individual stays in dependent employment. More interestingly, our study points at differences in the returns to self–employment that depend on the *type* of self–employment experience. Specifically, we find that negative effects on later wages in dependent employment only arise if the self–employment experience is from a sector different from the subsequent wage–employment sector. Furthermore, individuals who hired workers or who enjoyed a high income during self–employment do not face adverse effects on their subsequent wages in dependent employment. Thus, it seems fair to conclude that although self–employment appears to be a bad experience on average it matters greatly if the formerly self–employed changed sectors after a completed self–employment spell or hired workers or enjoyed a high income during his self–employment spell.

The rest of the paper is structured as follows. In Section 2 we present our theoretical framework, in Section 3 we describe our data. Section 4 outlines the empirical approach, while Section 5 contains the results. Section 6 concludes.

2 Theoretical framework

The basic premise of human capital theory, as developed by Becker (1964), is that wages are influenced by human capital which in turn is influenced by investments in education and training by individuals throughout their lives. Later, Mincer (1974) and others have pointed out that human capital can be acquired both in school and through labor market experience. Subsequently, a vast amount of empirical papers has confirmed this by estimating Mincer–type wage equations, finding substantial wage returns to both formal schooling and labor market experience.

The basic question we ask in this paper is whether the experience acquired in self–employment produces more or less human capital than experience acquired in wage–employment. As the tasks undertaken by the self–employed often differ from those performed by the wage–employed (see, e.g., Lazear, 2005), and because wage–employed and self–employed individuals may participate in different training activities, it seems natural to expect different effects. Spence (1973) argues that individuals educate themselves not to improve their human capital but merely to signal their innate abilities to future employers. The idea is that more able individuals can more easily (more cheaply) acquire education and therefore use it as a signal of ability. Following that line of thought, experience from wage–employment and self–employment may not only provide different amounts of human capital, but may also provide different signals about innate abilities to prospective employers.

In any case, whether self–employment experience builds human capital and/or provides a signal to future employers, it may be hypothesized, or even expected, to have different effects on subsequent wages than experience from wage–employment. This is what we set out to empirically test in the present paper.

In the literature, a distinction is typically being made between general and firm–specific human capital. Actually, Becker (1964) already distinguished between general and specific training in his original contribution. The latter is only valuable in the current job while the former has general applicability. This distinction can be used to explain why most empirical studies find substantial returns to firm–specific experience, also called "tenure", in excess of the returns to general labor market experience.

Lazear (2003) argues that firm–specific human capital really is about the right mix of general skills. That is, firms need different combinations of general skills, and through employment in a given firm, individuals acquire the right mix of skills over time. This mix, however, is no longer the right one when the individual changes jobs. This results in a lower value of the worker's human capital in the new job and hence a lower wage.

This line of reasoning suggests that the sectoral transition pattern in connection with the self–employment experience may also matter for subsequent wage effects. Experience from the same industry might provide human capital (and possibly also signals) which are more useful (more valuable) for subsequent employers, as it will result in a more appropriate mix of skills. If a period of self–employment is associated with a change of industry, it is thus important to analyze whether the resulting effects are due to the self–employment experience or due to the change in the sector of employment — or both. We therefore analyze the importance of the industry–specific experience when comparing the effects of wage–employment and self–employment experience on wages.

Finally, according to signaling theory, a spell of self–employment that involves employees or high income could provide a stronger signal of innate abilities than a spell without these characteristics. Thus, we also analyze the importance of such characteristics of the self–employment spell for the effects of self–employment experience on subsequent wage–employment wages.

3 Data

We use data from the Integrated Data Base for Labor Market Research ("IDA") compiled by Statistics Denmark. IDA contains register data on all individuals with Danish residence. The data base provides detailed information on experience in different occupations, hourly wages and a wide range of other individual–specific characteristics like educational background.[1]

The occupation of an individual in a given year is determined by Statistics Denmark according to the individual's primary labor market status in the last week of November. In the following, we shall distinguish between self–employment, wage–employment, non–employment and unemployment. A self–employed is defined by Statistics Denmark as an individual who owns a non–incorporated business, and either has employees or enjoys a wage–employment income below a fixed threshold, and is not registered as unemployed or non–employed. The common feature of our treatment group individuals hence is that their primary outlet for their labor is their own non–incorporated business.

The fact that Statistics Denmark records occupations only once a year means that we are unable to control for flows between labor market statuses within a year. We may for example miss some short spells of self–employment, non–employment or unemployment for those who are recorded as being wage–employed. Our results should therefore be interpreted as applying to self–employment spells exceeding a certain magnitude. However, if individuals with short spells are mis–categorized as consecutively wage–employed and short spells are not inherently different from longer spells, this

[1] For more details on the IDA data; see Abowd and Kramarz (1999).

may serve to diminish any difference between the treated and the non–treated – a sort of attenuation bias.

In the analysis that follows, we consider individuals who were full–time wage–employed in both 1990 and 1996 in the non–primary sector. We restrict attention to males aged 31–59 years in 1996 to avoid problems with early retirement and to make sure that all individuals have finished their educations. We discard individuals that encountered spells of unemployment or non–employment because we compare them to *consecutively* wage–employed individuals who by definition were not unemployed or non–employed either. I.e., we do not attempt to draw any inference about formerly self–employed who also encountered spells of non–employment or unemployment.

Furthermore, we drop immigrants from our data since some variables, most importantly the educational ones, are poorly measured for this group. We also delete the upper and lower percentile of the wage distribution as these extreme observations are potentially mis–reportings.

Table 1 displays descriptive statistics of occupations and transitions between wage–employment and self–employment. The upper part of the table shows the distribution of occupations in 1990 and 1996, respectively. The share of self–employed individuals was 7.5% in 1990 and 8.2% in 1996.

The middle part of Table 1 shows that 88.2% of the individuals who were wage–employed in 1990 were also wage–employed in 1996, whereas 3.3% were self–employed. Furthermore, 11,178 of the individuals who were wage–employed in both 1990 and 1996 experienced at least one spell of self–employment of at least one year. The average self–employment experience was 1.8 years with a median of one year.

The lower part of Table 1 further illustrates the duration of self–employment spells for workers being wage–employed in 1990. Most self–employment spells were relatively short. Of those individuals who became self–employed in 1991, 27.2% were back in wage–employment the following year; only 35.6% were still self–employed in 1996.

[Table 1: Occupations and Transitions between Wage–Employment and Self–Employment]

Table 2 displays descriptive statistics of the characteristics of the 10,436 individuals who experienced a spell of self–employment between 1990 and 1996 and for whom we observe the full history of sector affiliations. The latter explains the slight difference in the number of observations compared to Table 1. First, distinguishing between seven different industries, 30.4% of the individuals were self–employed in the same sector as they were later wage–employed in, and 21.3% were self–employed and wage–employed in the same sector in both 1990 and 1996. Second, relatively few spells of self–employment were followed by spells of unemployment or non–employment, as shown in the bottom part of Table 2. This might indicate that few exits from self–employment were forced quits.

[Table 2: Characteristics of Individuals with Self–Employment Experience]

Note that the Danish labor market is characterized by a high degree of flexibility, as firing costs are extremely low. In this respect, Denmark — and hence probably also our estimation results — compares better to the US and UK labor markets than to that of continental Europe. At the same time, however, the Danish welfare state takes care of the unemployed with particularly high compensation rates which is why the Danish model is often termed "Flexicurity" – a combination of flexibility in terms of hiring and firing opportunities and social security (Andersen and Svarer, 2006).

4 Empirical approach

4.1 Propensity score matching

Our aim is to estimate the average effect of the treatment (the self–employment spell) on the treated (ATT). The outcome variable is the natural logarithm of an individual's hourly wage rate in 1996 which means that our estimated ATTs translates into the relative effects of a previous self–employment spell on 1996 wages in wage–employment.

The "basic" treatment we consider is a spell of self–employment between 1990 and 1996 for individuals who were full–time wage employed in both 1990 and 1996. Furthermore, we distinguish between several different types of treatments as explained in Section 4.6 below.

We use propensity score matching (PSM) to address the potential sample selection problem. The idea behind PSM is to find a control group individual for each treatment group individual that has the same, or at least very similar, observed characteristics.

PSM is very well described in the literature, e.g., by Becker and Ichino (2002), Caliendo and Kopeinig (2008), Heckman et al. (1997) and Lechner (2001). The seminal references are Rosenbaum and Rubin (1983, 1985). A recent paper that applies PSM on a ten percent subsample of our data is Simonsen and Skipper(2006). Our exposition of PSM follows Caliendo and Kopeinig (2008).

In a natural experiment, a fraction of individuals would be randomly allocated to a (short) spell of self–employment before being returned to wage employment, while the remaining individuals would be kept in wage–employment throughout the period. The mean difference in the observed subsequent wages in wage–employment between these two groups would then constitute the ATT.

In practice, however, we are not able to allocate individuals randomly into occupations. PSM is an approach to mimic natural experiments (Rubin 1974; Rosenbaum and Rubin, 1983, 1985). It tries to find "clones" in terms of observable characteristics in the control group for each individual in the treatment group. The selection into the effective control group is not non–random but determined by how close each individual from the potential control group is to the individuals in the treatment group. In this way, PSM tries to eliminate any systematic differences between the treated individuals and the control group individuals. PSM is described in more detail in the following subsection. The validity of PSM hinges on the assumption that we are able to eliminate all systematic differences affecting both the outcome (the resulting wage) and the selection into the treatment (the self–employment experience).

4.2 The Average Treatment Effect on the Treated

Our parameter of interest is the ATT which is defined as

$$\tau_{ATT} = E[Y(1)|D = 1] - E[Y(0)|D = 1], \tag{1}$$

where $Y(1)$ denotes the outcome variable (the hourly wage) for treated individuals and $Y(0)$ denotes the outcome for untreated individuals. The counterfactual average for the treated individuals, $E[Y(0)|D = 1]$ is unobserved which forces us to replace it by an estimate. Approximating $E[Y(0)|D = 1]$ with the mean outcome of the untreated, $E[Y(0)|D = 0]$, is likely to generate biased estimates since variables that determine selection into treatment will also affect the outcome variable, leading to a "self–selection bias" which is defined as

$$\text{Bias} = \tau_{ATT} - E[Y(1)|D = 1] - E[Y(0)|D = 0]. \tag{2}$$

The purpose of econometric matching methods is to minimize the quantity $E[Y(1)|D = 1] - E[Y(0)|D = 0]$. Our approach to minimize it is to assume that, given a set of observable characteristics, $\boldsymbol{x}$, — which are not affected by treatment — potential outcomes are independent of the assignment to treatment. This assumes that selection is based on observable characteristics only and that all variables that influence treatment assignment and potential outcomes are simultaneously elements of $\boldsymbol{x}$.

This assumption is called "conditional independence" or "unconfoundedness". While we do control for a large set of relevant variables that are known to affect both wages and selection, we cannot formally test if the conditional independence assumption is indeed satisfied. We do, however, formally test whether treatment and control observations indeed no longer differ significantly with respect to observable characteristics after matching. This is known as the "balancing property".

As long as the set of observed characteristics $\boldsymbol{x}$ is small, one can exactly match treatment and control group observations to one another based on the individual elements of $\boldsymbol{x}$. Matching on a large set of observed characteristics, as we do in this paper, makes one–to–one matching infeasible. In order to overcome this dimensionality problem, Rosenbaum and Rubin (1983) suggest to use "balancing scores". They show that if potential outcomes are independent of treatment conditional on observed characteristics (which we assume), they are also independent of treatment conditional on a balancing score, $b(\boldsymbol{x})$. A frequently used estimator for the balancing score is the propensity score, the probability of treatment conditional on $\boldsymbol{x}$, a quantity that we estimate by a binary probit model.

An additional condition for our identification strategy to hold is the "common support" requirement. It rules out that the probability of treatment is perfectly predicted by $\boldsymbol{x}$. It makes sure that individuals with the same observed characteristics have a positive probability of receiving both treatment and non–treatment (Heckman et al., 1999).

The PSM estimator for the ATT is

$$\hat{\tau}_{ATT}^{PSM} = E_{\hat{P}(\boldsymbol{x})|D=1}\{E[Y(1)|D=1,\hat{P}(\boldsymbol{x})] - E[Y(0)|D=0,\hat{P}(\boldsymbol{x})]\}, \tag{3}$$

where $\hat{P}(\boldsymbol{x}) = \hat{b}(\boldsymbol{x})$ denotes the predicted probability of treatment. Equation (3) shows that the PSM estimate of the ATT simply is the mean difference in outcomes over the common support.

4.3 Matching method

Econometricians have developed a battery of methods to map treatment and control observations based on the propensity score as reviewed by Caliendo and Kopeinig (2008). These methods trade estimation bias against estimation precision (variance). We have tried the following methods: one–to–one matching without replacement; one–to–one matching with replacement; nearest neighbor matching with 2, 5 and 20 neighbors; kernel matching with an

Epanechnikov kernel with bandwidths 0.06, 0.18 and 0.54; kernel matching with a normal kernel with bandwidths 0.06, 0.18 and 0.54. Nearest neighbor matching is not a preferred option in the present setting since it tends to produce less unbiased results than the other methods. It is mostly applied in cases where there are only few control group observations available and we use it only to check the extent to which our results depend on the matching method. The estimated ATTs differ only slightly between the alternative matching models.[2] Standard errors are slightly larger for one–to–one matching and for kernel matching compared to nearest neighbor matching. We finally opted for one–to–one matching with replacement since it reduces estimation bias at the cost of higher variance (our standard errors are hence more conservative) and since it is computationally less burdensome than kernel matching while producing almost identical point estimates and standard errors.

4.4 Conditioning variables

To compute the propensity score, we first estimate a probit model for selection into treatment. As explained above we should condition on all variables that affect both treatment assignment *and* outcomes. Such a set of conditioning variables would ideally include broad range of skills/abilities measures, initial employment prospects and preferences for risk and job types. While controlling completely for all this is, of course, impossible, we try to proxy these variables by the set of explanatory variables described in the next few paragraphs. These variables prove to statistically and economically significantly affect both selection into self–employment and wage-employment wages, and as we explain below, they are likely to provide good controls for skills, employment prospects and preferences.

We use the following set of explanatory variables that prove to statistically and economically significantly affect both selection into self–employment and wage–employment wages: (i) years of tenure and years of tenure squared; (ii) age and age squared; (iii) sector affiliation: dummy variables for being employed in manufacturing, electricity, construction, trade, transport or finance with being employed in the public sector as the reference sector; (iv) education: dummy variables for twelve, 14, 16 and 18 years of education with

[2]The largest difference between the alternative estimators is 16% (for our "main" treatment/control combination that is discussed in Subsection 4.6 below). This is the case for a comparison between nearest neighbor matching with 20 neighbors and one one–to–one matching without replacement. The related difference for nearest neighbor matching with two neighbors is ten percent. Nearest neighbor matching generates larger point estimates than one–to–one matching in both cases.

less than twelve years of education constituting the reference category; (v) marital status; (vi) dummy variables for the number of children being one, two, or more than two with having no children being the reference category; (vii) dummy variables for living in a rural or an urban area with living in the Greater Copenhagen area as the reference category; (viii) the regional unemployment rate; (ix) the natural logarithm of the number of employees of the employer in 1990 and its square and (x) the hourly wage received in the initial employment in 1990;

Our selection of covariates follows Bruce and Schuetze (2004) with the exception of the firm size variable. All variables refer to 1990 in order to ensure that our conditioning vector $\boldsymbol{x}$ is not affected by treatment. If we used 1996 values instead, these variables would have be partly determined by a previous spell of self–employment, thereby violating the conditional independence assumption.

Tenure and age are both variables that are very likely to affect both wages and selection — and they prove to do so in our study as well. It is important to notice that by matching on tenure, we effectively take into account cohort effects as we compare treated and control group observations with the same length of tenure in 1990. Studies dealing with the internal economics of the firm like Baker et al. (1994a,b) as well as Gibbons and Waldman (1999, 2006) show that wage rates differ across entry cohorts independent of the characteristics of the individuals from different cohorts and that these wage differences are persisting.

That education is a major ingredient of any wage regression has been known at least since Mincer (1974). It is also likely to affect occupational choice as shown by Lazear (2005) and Lucas (1978).

Marital status and the number of children may be considered as proxy variables for risk aversion as discussed by, e.g., Leigh (1986), Halek and Eisenhauer (2001) as well as Schooley and Worden (1996). At the same time, marital status and the number of children also affects wages as shown by, e.g., Hamilton (2000).

An individual's geographical location is also likely to affect both selection and wages since employment prospects tend to be more gloomy in more rural areas. The regional unemployment rate is included since high unemployment rates may push some individuals into self–employment (Pfeiffer and Reize, 2000) and also affect wage rates (Akerlof and Yellen, 1990).

We include the number of employees at the establishment in 1990 as an explanatory variable since larger firms tend to pay higher wages (Brown and Medoff, 1989; Winter–Ebmer and Zweimüller, 1999). We are able to control for establishment size since our register data allows us to combine individual–level data with establishment–level data (matched employer–employee data).

While a positive relationship between current establishment size and wages is well–established in the literature, it is not so obvious that establishment size in 1990 affects current wages. Establishment sizes in 1990 and 1996 are, however, highly correlated which implies that individuals who were employed with a large (small) employer tend to prefer to be employed with a large (small) employer later on as well, or that larger firms tend to select the better workers in which case the establishment size acts as a proxy for unobservable individual differences.

Furthermore, it also turned out that establishment size has a statistically highly significant effect on selection as discussed in Subsection 5.1.

Finally, "initial" wages are used by Williams (2000) as well as Bruce and Schuetze (2004) as proxy variables for unobserved individual–specific heterogeneity.

Clearly, our set of conditioning variables only controls for *observed* individual characteristics while the "ideal" set of conditioning variable would include more direct measures of e.g. initial employment prospects and risk preferences.

Table 3 displays descriptive statistics for the variables involved in the estimations. It differentiates between treatment and control observations and also contains tests for significant differences in the variables used in the estimations.[3]

[Table 3: Summary Statistics]

The table shows that formerly self–employed individuals earn on average (and uncontrolled for observed characteristics) four percent more than consecutively wage–employed individuals, a difference that also is statistically highly significant. The difference in initial 1990 wages is seven percentage points but statistically insignificant. Individuals in the treatment group have fewer years of tenure, are less likely to work in manufacturing, electricity and transport, are less likely to work in trade and finance, are more likely have 18 or more years of schooling and are employed in smaller establishments.

There are no statistically significant differences with respect to age, marital status, the number of children, the geographical location, and years of education other than 18 or more years.

The aim of propensity score matching is to remove any differences in the distribution of these observed characteristics. If these are removed, the "balancing property" is satisfied.

[3]The difference in the number of observations in Table 3 and Table 2 is due to missing values for some variables used in the estimation.

4.5 Identification strategies of existing studies

While our approach is to find control individuals that resemble our group of previously self–employed individuals most closely, the existing studies briefly reviewed in the introductory section rely on OLS estimation that controls for a set of covariates that are likely to affect wages. Williams (2000) as well as Bruce and Schuetze (2004) use initial wages as a main control variable.

Identification is still, however, based on the assumption that selection into self–employment is random conditional on the covariates controlled for. There is no weighting by how close treatment and control observations are in terms of their observed characteristics and all effects are assumed to be linear. In Subsection 5.5 we briefly discuss alternative estimation results we produced based on the identification strategy used by existing studies.

Moreover, existing studies also uses current instead of before treatment covariates which is inappropriate since current values of covariates are likely to be determined by treatment, thus violating the Conditional Independence Assumption (which also needs to hold for OLS estimation). Treatment individuals will for example by definition have fewer years of tenure than control group individuals. This may conceal the effects of self–employment spells.

4.6 Treatment and Control Groups

Following the theory outlined in Section 2, our interest is not only in the general effect of a self–employment spell (our "basic" treatment) on subsequent wages, but also in the effects of different types of self–employment spells (our specialized treatments). We define a total of seven different treatments below. To provide appropriate counterfactuals for the different treatments, we also define six different control groups which mirror the different treatment groups.

Definitions of treatment groups

The "basic" treatment group (T_1) we consider consists of individuals with at least one spell of self–employment in the years 1991–1995. This treatment is similar to the one previously considered in the literature; see, e.g., Bruce and Schuetze (2004). In our estimation sample, 8,006 individuals receive this treatment.

We refine treatment group T_1 by (i) taking into account sector switching (treatment groups T_2 and T_3) and (ii) considering having had employees as self–employed or having enjoyed a particularly high income when self–employed (treatment group T_4).[4]

[4] Our measure of self–employment income is, just as in Hamilton (2000), net profits. A "high income" in self–employment is defined as an income in the fourth quantile in the

We distinguish two different overall transition patterns among the individuals in T_1: those who return to their old sector of wage–employment (5,472 individuals) and those who move to a different sector of wage–employment (3,910 individuals). It seems reasonable to expect that the effects on subsequent wages are different in the two cases. In the first case, workers are bringing back experience to their old sector. Workers may even know from the beginning that their self–employment spell is going to be temporary. In other words, it may be something that workers do to improve their subsequent wage in dependent employment in that sector. In the second case, workers proceed to a different sector following a self–employment episode.

For these reasons, we split up the basic treatment group T_1 into individuals who proceed to the same sector of wage–employment, treatment T_2, and those who move to a different sector, treatment T_3.

In the case of individuals who stay in the same sector of initial wage–employment, T_2, we distinguish between individuals who obtain self–employment experience from the *same* sector, treatment T_{2a}, and those who obtain it from a *different* sector, treatment T_{2b}.[5] A priori, we expect the former type of experience, T_{2a}, to be valued more by subsequent employers than the latter type of experience, T_{2b}.

Among the individuals in treatment group T_3 — those whose sector of employment differs between 1990 and 1996 — we also distinguish between individuals who were self–employed in the *same* sector as the one in which they subsequently become wage–employed, treatment T_{3a}, and those who were self–employed in a *different* sector, treatment T_{3b}. We again expect the former type of experience to be most valuable in subsequent wage work.

Clearly, studying wage effects of sector switching implies that we are potentially introducing an additional selection problem. Note, however, that we compare self–employed individuals who switched sectors to consecutively wage–employed individuals who also switched sectors. I.e., we compare individuals that followed the same employment patterns. Hence, we may be unable to identify absolute effects on wages but we do properly identify parameter of our main interest, the difference in wage rates for treatment and control group observations who both changed sectors.

Finally, we consider two treatments which may provide valuable signals to future employers or which may contain particularly valuable experiences. Treatment T_{4a} consists of individuals with at least one employee in one of the years as self–employed.[6] We additionally consider individuals with earnings

self–employment income distribution.

[5] When individuals have self–employment experience from more than one sector in the years 1991 to 1995, we use the sector of the last year of self–employment.

[6] The vast majority of former entrepreneurs had at most one employee which is why

in self–employment in the upper quartile among the self–employed, treatment T_{4b}.[7] Clearly, self–employment income will depend on taxation and it is certainly possible that some individuals have high de facto but low taxable incomes (the variable we measure). Similarly, the number of employees in self–employment may reflect low risk–aversion.

Definitions of control groups

Our most general group (C_0) consists of individuals who were permanently wage–employed between 1990 and 1996. This is the central control group considered by Bruce and Schuetze (2004).

The focus of our attention is, however, on control group C_1 which consists of individuals who changed jobs at least once in the period. This is our "basic" control group since our treated individuals — by definition — also changed their job between 1991 to 1995.

Furthermore, we split our basic control group into individuals who stayed in the *same* sector in 1990 and 1996, C_2, and those who move to a *different* sectors, C_3. We proceed that way in order to provide more relevant counterfactuals for treated individuals who return to their old sector, e.g. treatment group T_2, and for those who join a different sector, i.e. treatment group T_3. Thus, treatment group T_2 compares to control group C_2 and treatment group T_3 compares to control group C_3.

Finally, to compare the effects of self–employment experience to those of unemployment and non–employment experience, we define our final two control groups as individuals who were wage–employed in 1990 and in 1996, but were unemployed (control group C_4) or non–employed (control group C_5) for at least one period in the years 1991 to 1995 while not having encountered a single spell of self–employment.

Figure 1 provides an overview of how we split up our treatment and control groups.

[Figure 1: Overview of Treatment and Control Groups Combinations]

Table 4 provides an overview of our treatments and controls — and the combinations of these considered in the following section. Our econometric

we only consider this type of treatment here, few had two and even fewer had more than two employees. In specifications not displayed in this paper, we additionally considered having had two or more than two employees as additional treatments. The corresponding estimation results indicate that these types of treatments go along with higher and statistically more significant positive returns to self–employment.

[7] We used the mean of the number of employees and mean self–employment income if individuals encountered multiple self–employment spells. Using maxima instead did neither qualitatively nor quantitatively alter the results.

approach is the same for all combinations of treatment and control groups. For the sake of brevity we do, however, only discuss details for our "basic" treatment/control group combination (T_1 and C_1) below.[8]

[Table 4: Treatment Group and Control Group Combinations]

Econometric software
We use the "psmatch2" module by Leuven and Sianesi (2003) implemented in STATA to perform the PSM estimation. The standard errors associated with the matching estimates in Table 5 are calculated according to Lechner (2001). Abadie and Imbens (2006) show that bootstrapped standard errors may be inconsistent in large samples when matching with replacement is used which is precisely our setting. Bootstrapped standard errors and Lechner's one do, however, differ very little in our estimations.

5 Results

We first discuss the estimation results for the binary probit model for selection into treatment since these form the basis for the propensity score and thus for the matching approach. We also provide tests for the quality of our matches. The PSM results for our various treatment and control group combinations are presented afterwards.

5.1 Binary probit results for selection

Main findings
A binary probit model is used to calculate the propensity score, $\hat{P}(\boldsymbol{x})$, the predicted probability of treatment. Estimation results are displayed in Appendix A which also contains the results of a more flexible empirical specification that we shall discuss below. They show that tenure has a negative effect on the probability of becoming self–employed. Age has an inverse U–shaped effect on selection with a maximum being reached at 45.7 years.

There is no statistically significant effect of marital status and the number of children. These "family status" variables are statistically jointly highly significant, however. Geographic location and regional unemployment rates are not found to having separately or jointly significant effects on selection.

We find a U–shaped relationship between establishment size in 1990 and the probability of becoming self–employed. That is, employees in very large

[8] Details for the other treatment/control group combinations are available from the authors upon request.

and very small establishments are more likely to become self–employed than employees in medium–sized establishments. The probability of becoming self–employed is, however, lowest for individuals who have been employed in an establishment with 2,166 employee which practically means that the propensity of becoming self–employed is decreasing in establishment size given a mean establishment size of 82.3 employees.

Of the set of education dummy variables, only the effect of 18 years of schooling or more is statistically significant (and negative).

Match quality

Since we do not condition on all covariates but only on the propensity score, we need to assess if our matching procedure is able to balance the distribution of treatment and control individuals, i.e., if the match quality is satisfactory.

Rosenbaum and Rubin (1985) suggest to use "standardized biases", which are simply the differences in the means of the covariates for treatment and control observations after matching, weighted by their standard deviations. It turns out that the parsimonious model does not match particularly well on the sector dummies. Caliendo and Kopeinig (2008) suggest interacting the variables that are not well matched with other covariates in an attempt to increase their importance in the matching process. We therefore interact the set of sector dummy variables with both initial wage and tenure. We also include the second to fifth polynomial of the initial wage since we did not match well on initial income after the inclusion of the sector dummy interaction. We use these newly constructed variables as additional covariates. Estimation results are shown in Appendix A (alongside the results of our initial specification without the interactions). Their interpretations are cumbersome, however, due to the many interactions. Hence, we do not discuss them here. The main result from this larger model is that there no longer exist statistically significant differences in the covariates between treatment and control group observations. This finding does not only apply to the combination of our "basic" control group and our "basic" treatment group but to the other combinations we consider as well.

Sianesi (2004) additionally suggests to re–estimate the propensity score on the matched sample, i.e., only including the treatment individuals and the matched control group observation, and compare the pseudo R^2's before and after matching. There should not exist a significant difference in the distribution of covariates after matching and the pseudo R^2's should therefore be close to 0. In addition, tests for joint significance of the covariates should reject joint significance after matching. The pseudo R^2's after matching are between 0 and 0.004 across our different treatment and control group combinations. Tests for joint significance easily reject the ability of the covariates to explain selection after matching as well.

Our formal tests hence indicate that we match treatment and control group observations very well. There is no indication of the balancing property not being satisfied. In addition, we test if the common support assumption is satisfied which is the case for all treatment/control combinations we consider.

5.2 ATT estimates

Table 5 displays the ATT estimation results for the alternative combinations of treatment and control groups using PSM. In the following subsections, we shall go through the results in detail, starting with the effects of the "basic" treatment (T_1, individuals with a spell of self–employment between 1990 and 1996) in Section 5.3. We discuss the more specific treatments in Section 5.4.

[Table 5: Effects of Self–Employment Experience on Subsequent Wages]

5.3 The "basic" treatment

The first row of Table 5 contains the estimated effects of the "basic" treatment, T_1. It shows that a spell of self–employment between 1990 and 1996 goes along with a reduction in hourly wages in subsequent wage–employment of 2.9%, an effect that is measured with relatively high precision given a standard error of 0.6%.

For comparison, Bruce and Schuetze (2004) find that an additional year of self–employment experience reduces the post self–employment wage in dependent employment by between 3.0% and 15.6% compared to one year of additional wage work experience. These effect are, however, often quite imprecisely estimated. Our results hence indicate somewhat less negative but statistically much more significant consequences of self–employment.

Bruce and Schuetze (2004) also perform a regression where they only use those individuals who also experienced a job change in wage–employment as the comparison group. This corresponds to our "basic" treatment/control combination T_1/C_1. This tends to make their results insignificant – possibly due to their substantially reduced sample size. Our results are more robust qualitatively to this change of control groups as we find a wage reduction of 1.6% (standard error 0.6%) once we compare our basic treatment group to control group individuals who changed jobs.

In sum, self–employment appears to be a bad experience compared to continued wage work since it leads to lower wage rates compared to consecutive dependent employment.

One reason for the negative effect of a self–employment spell could be that formerly self–employed become employed by a smaller firm in 1996 compared to 1990. Given that smaller firms pay less than larger firms, this would have adverse effects on wages. To analyze this issue, we calculate the differences in employer sizes between 1990 and 1996 for consecutively wage–employed individuals who changed jobs at least once — our "basic" control group T_1 — and individuals who encountered at least one spell of self–employment — our "basic" treatment group C_1. It turns out that there is no evidence for formerly self–employed workers are more likely to join smaller firms than consecutively wage–employed individuals: there is no statistically significant difference in the mean and the median change in employer size between 1990 and 1996 between our treatment group individuals and our control group individuals.

However, formerly self–employed individuals fare better than formerly unemployed or non–employed individuals — individuals in control group C_4 and C_5, respectively. A comparison of our basic treatment group to these two control groups indicates positive wage differences of around 15% that are measured with high precision. For comparison, Bruce and Schuetze (2004) estimate that an additional year of unemployment is associated with a wage reduction of between 8.2% and 55.7% relative to a year of continued wage–employment.

Another difference of our results compared to those of Bruce and Schuetze (2004) is that their estimates become smaller in magnitude (and less significant) when they control for initial wages. They find that the less able individuals tend to encounter self–employment spells. We find the opposite phenomenon in our OLS regressions that we briefly discuss in Subsection 5.5. This may reflect a higher degree of wage compression in the Danish labor market. Wage compression increases the incentives for more able individuals to engage in self–employment as shown by Malchow–Møller et al. (forthcoming).

5.4 Specific Treatments

Now that we have estimated the effects of our basic treatment, T_1, on subsequent wage–employment wages, we split up the treated individuals into those who return to their old sector of wage–employment, T_2, and those who move on to a different sector, T_3.

Same sector of wage–employment in 1990 and 1996

We consider first formerly self–employed individuals who in 1996 return to their initial 1990 sector of wage–employment, T_2. In this case, we do not find statistically significant effects of the treatment when the counterfactual used is individuals who changed jobs between 1990 and 1996 (control group C_1). If we instead compare treatment group T_2 individuals to control group individuals who were employed in the same sector in 1990 and 1996 (control group C_2), we find a negative wage effect of 2.6% of treatment T_2 relative to counterfactual C_2. Taking these two results together implies that formerly self–employed individuals do not encounter wage differences relative to consecutively wage–employed with a job change, but that there are significant wage differences once they are compared to the more narrowly defined group of control group job changers who do not switch sectors of employment. This is a first indication of the importance of sectoral transitions changes for subsequent wages.

We expect the treatment effect to be more positive (or less negative) if the self–employment experience is from the same sector as the subsequent wage–employment sector. In order to quantify these effects, we split the T_2 treatment into self–employment experience from the same sector, T_{2a}, and self–employment experience from a different sector, T_{2b}. As expected, only the latter treatment is associated with negative and statistically significant effects when compared to the C_2 counterfactual. The point estimate for this ATT is -3.7%. By contrast, formerly self–employed individuals who become wage–employed in their previous sector of self–employment do not encounter a statistically significant wage difference when compared to the C_2 counterfactual.

Different sector of wage–employment in 1990 and 1996

We now consider those individuals whose wage–employment sector in 1996 is different from their wage–employment sector in 1990. This is the T_3 treatment. Compared to our basic counterfactual C_1 we find a significantly negative effect of this type of treatment of 2.3%.

However, if we compare the T_3 individuals to individuals in wage work who also changed sectors between 1990 and 1996 (control group C_3), the negative effect vanishes. This suggests that changing sectors after a completed self–employment spell is a main driver of the generally negative effect of past self–employment spells on subsequent wages found by our "basic" treatment/control specification discussed in Subsection 5.3.

To investigate this issue even further, we split the T_3 treatment group up into individuals who do not leave their self–employment sector after their self–employment spell ended, treatment group T_{3a}, and those individuals who did change sectors after the self–employment spell, treatment group T_{3b}. When compared to the C_1 control group, the effect is -4.1% for the latter group. The significance of this effect disappears when treatment group T_{3b} is compared to consecutively wage–employed individuals who changed sectors between 1990 and 1996 (control group C_3). For the former group, those who obtain self–employment experience in the same sector as the one in which they are subsequently wage–employed, there is a positive effect of around six percent — both when compared to C_1 and to C_3. However, we should be aware of that treatment group T_{3a} contains 682 individuals only.

In sum, as for those who return to the same sector of wage–employment, T_2, negative effects of a self–employment spell can only be found when the self–employment spell is associated with a subsequent change of sectors.

Finally, Table 5 shows that formerly self–employed individuals who hired workers, T_{4a}, or enjoyed a high income while self–employed, T_{4b}, receive at least the same wage rate as consecutively wage–employed individuals who changed jobs, C_1. In the latter case, the estimated effect of 6.3% is in fact significantly positive. This indicates that these types of self–employment experiences provide either more valuable signals to future employers or contain more valuable experiences than the average self–employment spell.

5.5 OLS findings

This section briefly discusses the results using OLS regressions since this is the approach taken by the existing studies briefly reviewed in Section 1. In this case, the ATT is the coefficient estimate on the dummy variable for treatment in an OLS regression of the explanatory variables on the natural logarithm of hourly wages. We use the same set of explanatory variables as for our propensity score matching model that includes polynomials and interactions between the explanatory variables (see Appendix A).

Appendix B displays OLS regression results where all explanatory variables are set to 1990 values to be comparable to our PSM results. There are both quantitative and qualitative differences between our PSM and OLS estimation results. The most striking qualitative difference is that OLS estimation generates a statistically highly significant and *positive* effect of the basic treatment, T_1, compared to the basic counterfactual, C_1. In cases where there are no qualitative differences between the OLS and the PSM estimation results, the PSM estimates tend to be larger in magnitude and to be statistically more significant.

Potential reasons for the differences between OLS and PSM estimation is the inability of OLS to balance and its linearity assumption. Our OLS regressions do contain polynomials and interaction terms already which also prove to be statistically highly significant. Adding additional interactions and high order polynomials does not lead to more similarity between OLS and PSM results which lets us conclude that the differences are due to OLS not being able to adequately balance treatment and control group observations.

6 Conclusion

This paper has estimated the effects of past self–employment experience on subsequent earnings in wage work using the population of Danish men between 31 and 59 years of age. We considered a cohort of individuals who were wage–employed in both 1990 and 1996 and who either were consecutively wage–employed within that period or who encountered a spell of self–employment. In order to deal with potential selection problems, we applied propensity score matching.

We found that a spell of self–employment is *in general* associated with a negative effect on subsequent hourly wages in dependent employment. If the treated individuals are compared to consecutively wage–employed individuals, we find that the former group receives an hourly wage that is 2.9% lower than that of the latter group. The effect is reduced to -1.6% once the formerly self–employed are compared to consecutively wage–employed individuals who changed jobs within the period considered. This counterfactual is perhaps more relevant since switching from wage–employment to self–employment (and back) involves a job change as well.

However, distinguishing between treatments that involve different sectoral transitions reveals that negative effects of self–employment spells are only found when the self–employment spell is followed by a sector change when the individual returns to dependent employment. Hence, the on average negative effect of a self–employment spell appears to be due to sector–switching rather then self–employment per se.

Moreover, self–employment spells with high income or employees are not found to be associated with negative wage effects which indicates that these spells provide more positive signals to or more valuable human capital for future employers.

Our findings thus underscore that there are divergent forces at work that determine the overall effects of self–employment experience on subsequent wages in dependent employment. In other words: self–employment is not always a bad experience.

References

Abadie, A., Imbens, G., 2006. Large sample properties of matching estimators for average treatment effects, Econometrica 74, 235–267.

Abowd, J.M., Kramarz, F., 1999. The analysis of labor markets using matched employer-employee data. In: O. Ashenfelter and D. Card (Eds.), Handbook of Labor Economics 3B, Elsevier Science, Amsterdam, 2629–2710.

Akerlof, G.A., Yellen, J.L., 1990. The fair wage–effort hypothesis and unemployment. The Quarterly Journal of Economics 105(2), 255–283

Andersen, T.M., Svarer, M., 2006. Flexicurity – den danska arbetsmarkadsmodellen. Ekonomisk Debatt 1, 17-29.

Baker, G., Gibbs, M., Holmström, B., 1994a. The internal economics of the firm: evidence from personnel data. Quarterly Journal of Economics 109(4), 881–919.

Baker, G., Gibbs, M., Holmström, B., 1994b. The wage policy of a firm. Quarterly Journal of Economics, 109(4), 921–955.

Becker, G., 1964. Human capital, National Bureau of Economic Research, New York.

Becker, S.O., Ichino, A., 2002. Estimation of average treatment effects based on propensity scores. The Stata Journal 2(4), 358-377.

Benz, M., Frey, B.S., 2004. The value of autonomy: evidence from the self–employed in 23 countries. University of Zurich, Institute for Empirical Research in Economics Working Paper No. 173.

Blanchflower, D.G., 2000. Self–employment in OECD countries. Labour Economics 7, 471–505.

Brown, C., Medoff, J., 1989. The employer size–wage effect. Journal of Political Economy 97(5), 1027–1059.

Bruce, D., Schuetze, H.J., 2004. The labor market consequences of experience in self–employment. Labor Economics 11, 575-598.

Caliendo, M., Kopeinig, S., 2008, Some practical guidance for the implementation of propensity score matching. Journal of Economic Surveys 22(1), 31–72.

Evans, D.S., Leighton, L.S., 1989. Some empirical aspects of entrepreneurship. American Economic Review 79, 519-535.

Ferber, M.A., Waldfogel, J., 1998. The long–term consequences of nontraditional employment. Monthly Labor Review 121 (May), 3-12.

Gibbons, R., Waldman, M., 1999. A theory of wage and promotion dynamics inside firms. The Quarterly Journal of Economics 114(4), 1321–1358.

Gibbons, R., Waldman, M., 2006. Enriching a theory of wage and promotion dynamics inside firms. Journal of Labor Economics 24(1), 59–107.

Martin, H., Eisenhauer, J.G., 2001. Demography of Risk Aversion. The Journal of Risk and Insurance 68(1), 1–24.

Hamilton, B., 2000. Does entrepreneurship pay? An empirical analysis of the returns to self–employment. Journal of Political Economy 108, 604–631.

Heckman, J., Ichimura, H., Todd, P., 1997. Matching as an econometric evaluation estimator: evidence from evaluating a job training programme. Review of Economic Studies 65(2), 261–294.

Heckman, J., LaLonde, R. , Smith, J., 1999. The economics and econometrics of active labor market programs. In: Ashenfelter, O., Card, D. (Eds), Handbook of Labor Economics 3, North Holland, Amsterdam, 1865–2086.

Lazear, E.P., 2003. Firm–specific human capital: a skill–weights approach, IZA Discussion Paper No. 813.

Lazear, E.P., 2005. Entrepreneurship. Journal of Labor Economics 23(4), 649–680.

Lucas, R., 1978. On the size distribution of business firms. Bell Journal of Economics 9(2), 508–523.

Lechner, M., 2001. Identification and estimation of causal effects of multiple treatments under the conditional independence assumption. In: Lechner, M. and F. Pfeiffer (Eds.), Econometric evaluation of labour market policies, Physica–Verlag, Heidelberg, 1–18.

Leigh, J.P., 1986. Accounting for tastes: correlates of risk and time preferences. Journal of Post Keynesian Economics 9, 17-31.

Leuven, E., Sianesi, B., 2003. PSMATCH2: Stata module to perform full Mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing, URL:
http://ideas.repec.org/c/boc/bocode/s432001.html.

OECD (2000), OECD employment outlook, Paris.

Malchow–Møller, N., Markusen, J.R., Skaksen, J.R., (forthcoming), Labour market institutions, learning and self–employment. forthcoming in Small Business Economics.

Mincer, J., 1974. Schooling, experience, and earnings. Columbia University Press, New York.

Niefert, M., 2006. The impact of self–employment experience on wages. Centre for European Economic Research mimeo, Mannheim.

Pfeiffer, F., Reize, F., 2000. Business start–ups by the unemployed – an econometric analysis based on firm data. Labour Economics 7, 629-663.

Rosenbaum, P., Rubin, D., 1983. The central role of the propensity score in observational studies for causal effects. Biometrika 70(1), 41–50.

Rosenbaum, P., Rubin, D., 1985. Constructing a control group using multivariate matched sampling methods that incorporate the propensity score. The American Statistician 39(1), 33–38.

Rubin, R., 1974. Estimating causal effects of treatments in randomized and nonrandomized studies, Journal of Educational Psychology 66, 688–701.

Schooley, D.K., Worden, D.D., 1996. Risk aversion measures: comparing attitudes and asset allocation. Financial Services Review S(2), 87–99.

Sianesi, B., 2004. An evaluation of the Swedish system of active labor market programs in the 1990s. The Review of Economics and Statistics 86(1), 133–155.

Spence, M., 1973. Job market signaling. Quarterly Journal of Economics 87, 355–374.

Simonsen, M., Skipper, L., 2006. The costs of motherhood: an analysis using matching estimators, Journal of Applied Econometrics 21(7), 919-934.

Williams, D.R., 2000. Consequences of self–employment for women and men in the United States. Labour Economics, 7, 665-687.

Williams, D.R., 2002. Returns to education and experience in self–employment: evidence from Germany, IRISS Working Paper No. 2002–04; www.ceps.lu/iriss.

Winter–Ebmer, R., Zweimüller, J., 1999. Firm–size wage differentials in Switzerland: evidence from job-changers. The American Economic Review Papers and Proceedings 89(2), 89–93.