Wage uncertainty and the labour supply of self-employed workers

Simon C. Parker*, Tim Barmby and Yacine Belghitar
Department of Economics & Finance
University of Durham

Abstract

We analyse the effects of wage uncertainty on the labour supply of self-employed workers, using PSID data on self-employed American males. The standard deviation of past wages, as a measure of wage uncertainty, is the key determinant of male self-employed labour supply, with a significant positive effect. In contrast there is no effect from the (instrumented) wage or other explanatory variables. Our findings are consistent with the self-employed 'self-insuring' in response to greater uncertainty by working longer hours, and they can also help explain why self-employed Americans work longer average hours for lower average wages than their employee counterparts.

JEL: D81, J22, J23

Keywords: labour supply, self-employed, uncertainty

Approximate word count: 6500

*Corresponding author. Contact address: 23–26 Old Elvet, Durham DH1 3HY, UK. E-mail: S.C.Parker@durham.ac.uk. Tel: +44 (0) 191 334 6351. This research was supported by ESRC Grant Number 000-22-0020. We are grateful to seminar participants at St. Andrews for helpful comments. Naturally, we are solely responsible for the interpretation of the data and any errors.

1 Introduction

The self-employed account for around 10 per cent of the workforce in most developed countries, employ a similar number of additional workers, and run most of these nations' firms (Parker, 2004). Despite this, relatively little is known about the labour supply of self-employed workers. In contrast, there is an extensive and growing literature on the labour supply of employees (Blundell and MaCurdy, 1999). If for no other reason, the burgeoning policy interest centring on the promotion of entrepreneurship and labour market flexibility (with which the self-employed are often associated) makes it timely to explore more closely what factors bear on the labour input of this important group of workers.

Although pronounced, the literature's neglect of self-employed labour supply has not been complete. Early work by Wales (1973) estimated a simple static model of work hours using a sample of American self-employed proprietors. More recently, researchers have explored the labour supply behaviour of specific self-employed occupations, including dentists (Boulier, 1979; Scheffer and Rossiter, 1983), physicians (Thornton, 1998), and taxi-drivers (Camerer *et al*, 1997). We depart from these studies by focusing on a factor that we argue is a particularly salient one for the self-employed: income risk. There is now clear evidence that self-employed workers have more variable and unequal incomes than employees do (Carrington *et al*, 1996; Parker, 1997). A separate body of evidence also shows that self-employed workers face a greater hazard of involuntary exit from their occupation (Evans and Leighton, 1989; Phillips and Kirchhoff, 1989; Taylor, 1999). That is a distinct form of risk that we do not focus on this paper. Instead, we focus specifically on uncertainty of returns *within* self-employment, as measured by the standard deviation of past wages, and ask how this affects the labour supply of self-employed workers.

After providing a theoretical discussion of the issues, we conduct an empirical analysis using a sample of PSID data on self-employed American working-age males. We exploit the panel nature of the PSID to construct our measure of wage uncertainty. We find this variable to be the key influence explaining male self-employed labour supply, with a significant positive effect. In contrast we find no effect from the (instrumented) current-period wage, once wage uncertainty is controlled for.

A positive effect of wage uncertainty on self-employed labour supply can potentially resolve a long-standing puzzle of why self-employed Americans work longer average hours for lower average wages than their employee counterparts do (Aronson, 1991; Carrington *et al*, 1996; Hamilton, 2000). With free occupational choice, it is unclear why this state of affairs should persist. Even if the self-employed accepted lower wages in return for a compensating differential such as love of independence, this could only explain lower wages, not longer work hours in self-employment. An al-

ternative explanation, that the self-employed have some kind of indivisibility in their production function, is also unsatisfactory, because it begs the question about what the indivisibility could be and why alternative technologies are not used instead.[1] And it does not seem convincing to argue that employees and the self-employed have substantially different tastes for leisure, or are located on different sections of a common labour supply schedule. Employees would need to be located on a backward bending section of a labour supply schedule in order for a wage premium to be associated with shorter working hours than the self-employed. But recent research findings cast doubt on the relevance of a backward bend in employee labour supply (Blundell and MaCurdy, 1999, Table 1).

Instead, we propose an alternative explanation based on a 'self-insurance' argument. If the self-employed respond to greater earnings uncertainty by working harder to make the deterministic part of their incomes larger, they might end up working longer hours than employees do, despite facing a lower hourly wage. And the existence of a valuable compensating differential in self-employment could be sufficient to establish this as an occupational equilibrium.

The remainder of the paper is structured as follows. Section 2 provides a theoretical analysis of the relationship between wage uncertainty and labour supply, and derives the conditions under which self-insurance can explain the above labour supply puzzle. Section 3 discusses empirical issues in estimating the relationship econometrically. Section 4 briefly describes the characteristics of our data sample. Section 5 presents the empirical results, and Section 6 concludes.

2 Theoretical considerations

Consider a set of identical self-employed individuals who work alone, and earn income

$$y = g(wL^* + B; \theta)\,, \tag{1}$$

where w is the hourly wage, L^* is the chosen number of hours worked (derived below), B is investment income, and θ is a random variable capturing uncertainty. The function $g(\cdot)$ can take several different forms, depending on whether the random shocks θ affect returns additively or multiplicatively, as we discuss below. Note that the wage w in (1) above might not be exogenous, but possibly depends on L^*, an

[1] The 'wage' of self-employees (and also of employees, in some data sets) is usually computed by dividing a measure of total income or profit by work hours (see, e.g., Hamilton, 2000). It might be thought that this automatically yields a negative relationship between the wage and work hours. However, quite apart from the fact that this argument does not bear on the employee-self-employed comparison, it is not generally correct, because it assumes that there is a fixed profit available in self-employment – a particular type of indivisibility.

issue we will return to later.

Suppose θ is distributed according to $F(\theta)$ and has a zero mean: i.e.,

$$\int_{-\infty}^{+\infty} dF(\theta) = 1 \qquad \textit{and} \qquad \int_{-\infty}^{+\infty} \theta \, dF(\theta) = 0 \, .$$

Assume further that all draws from the random variable are independent and identically distributed (i.i.d.). This assumption represents the idea of an idiosyncratic shock that affects incomes in such a way that all risk-averse people would prefer to be fairly insured against such fluctuations. It is assumed for simplicity that while firms can offer insurance against these fluctuations to every employee (as in, e.g., Schaffner, 1993), the self-employed are unable to purchase actuarially fair insurance. For example, this might be because moral hazard problems are too severe to enable such a market to exist, or because few self-employed people run enterprises with sufficiently high turnover to diversify their risks through a stock market (Kihlstrom and Laffont, 1983).

For simplicity, utility is assumed to be both concave and separable. Separability merely facilitates easy exposition of the key results; we will signpost relaxations of this assumption below. Let $\mathcal{E}$ denote expected value, $\mathcal{V}$ variance, and $\mathcal{C}$ covariance. Then expected utility is

$$\mathcal{E}[u(y, L)] = \int_{-\infty}^{+\infty} u(y) \, dF(\theta) + v(L)$$

where $u'(\cdot) > 0$, $u''(\cdot) < 0$, $v'(\cdot) < 0$, $v''(\cdot) < 0$, and all of these functions are continuous with well behaved second and third derivatives.

To derive our results on self-employed labour supply, we propose a simple model of self-employment income generation that leads to two different specifications of (1). Reflecting common practice, the self-employed wage w is defined as the ratio of earnings, E, to work hours, L. But net reported self-employed earnings typically contain not only the direct return to labour – as is the case for employees – but also an autonomous component, such as high inventory sales in times of high demand (on the positive side) and lump-sum business operating costs (on the negative side). Thus we write $E = E_0 + \alpha L$, where $\alpha > 0$ is the marginal product of labour and where E_0 is autonomous. Not only is this specification an arguably realistic representation of income for the self-employed, but also it has the advantage of presenting a convenient form within which to analyse two different kinds of stochastic shocks: additive and multiplicative. These are analysed in the next two subsections. At the end of each, we shall consider the extent to which risk bears on the labour supply puzzle identified in the Introduction.

2.1 Additive shocks

We first consider the case where E_0 is random (with mean $\overline{E}_0$) but α is fixed, i.e.,

$$g(wL^* + B; \theta) = (\overline{E}_0 + \theta) + \alpha L^* + B\,. \tag{2}$$

Hence in this specification the computed wage is risky in the sense that $w = \alpha + (\overline{E}_0 + \theta)/L^*$ is stochastic; while (2) shows that shocks affect incomes additively. This case might arise, for example, if demand shocks determine how much inventory is sold in any given time period. The first order condition for self-employed labour supply is

$$\alpha\mathcal{E}[u'(\cdot)] + v'(\cdot) = 0\,,$$

yielding the labour supply function $L^* = L^*(\alpha, B)$, which is assumed to be non-decreasing in α.

Using the results of Rothschild and Stiglitz (1971), a mean preserving spread (MPS) in θ increases or decreases L^* as $A(\theta) := \alpha u'(\cdot) + v'(\cdot)$ is convex or concave in θ, respectively. Differentiating twice with respect to θ yields

$$A''(\theta) = \alpha u'''(\cdot)\,. \tag{3}$$

The right hand side of (3) is positive if $u'''(\cdot) > 0$, in which case $A(\theta)$ is convex, and labour supply increases with the level of uncertainty. As Arrow (1971) and others (e.g., Kraus and Litzenberger, 1976) have maintained, $u'''(\cdot) > 0$ is a property possessed by reasonable utility functions, including those exhibiting nonincreasing absolute risk aversion. Hence this 'usual' case has the self-employed 'self-insuring' in response to greater uncertainty, choosing a larger labour supply and thereby making the deterministic part of their income relatively larger.[2]

The interest of this result is underscored by the fact that it can potentially resolve the labour supply puzzle. Suppose there is free occupational choice between self-employment and paid employment, the latter denoted by PE superscripts; that there is a non-pecuniary compensating differential in self-employment, whose utility value is given by ϕ; and that, for simplicity, there is no uncertainty in paid employment.[3] In equilibrium, the paid employment wage w^{PE} adjusts to ensure that all individuals

[2]If utility is quadratic, $u'''(\cdot) = 0$. Then self-employed labour supply is invariant to the level of risk. If utility is non-separable, then convexity of $A(\theta)$ also requires that the third cross-partial $\partial^3 u/[\partial L^*(\partial y)^2]$ is positive.

[3]The assumption of no wage uncertainty in paid employment is invoked only for expositional ease. The argument still goes through if there is uncertainty in paid employment, as long as there is greater uncertainty in self-employment.

are indifferent between the two occupations:

$$u(w^{PE}L^{*PE} + B) + v(L^{*PE}) = \mathcal{E}[u(\overline{E}_0 + \theta + \alpha L^* + B] + v(L^*) + \phi\,, \tag{4}$$

where L^{*PE} is optimal labour supply in paid employment, derived by differentiating the LHS of (4) and setting to zero. The stylised facts (Aronson, 1991; Carrington *et al*, 1996; Hamilton, 2000) are that (i) $L^* > L^{*PE}$ and (ii) $w^{PE} > w$. This jointly implies that, for (4) to hold, we must have $\phi > 0$. Then as we saw above, the (greater) wage uncertainty in self-employment can explain the greater labour supply of the self-employed despite labour supply schedules being non-downward sloping in own wage and the average self-employed wage being less than that of employees. That is, it is generally possible to find a ϕ value such that (4) holds at the same time as both $w < w^{PE}$ and $L^* > L^{*PE}$.

2.2 Multiplicative shocks

Next we consider the case where E_0 is fixed but the marginal product of labour is random (with mean $\overline{\alpha}$), i.e.,

$$g(wL^* + B; \theta) = E_0 + (\overline{\alpha} + \theta)L^* + B\,. \tag{5}$$

Hence in this specification the wage is risky in the sense that $w = (\overline{\alpha} + \theta) + (E_0/L^*)$ is stochastic; while (5) shows that shocks affect incomes multiplicatively. This case might arise, for example, if shocks affect output prices. The first order condition for self-employed labour supply becomes

$$\mathcal{E}[(\overline{\alpha} + \theta)u'(\cdot)] + v'(\cdot) = 0\,. \tag{6}$$

From this it is possible to derive optimal labour supply $L^* = L^*(\overline{\alpha}, B)$, which is assumed to be non-decreasing in $\overline{\alpha}$ as before.

The comparative statics are somewhat more involved than for the additive case. Writing $A(\theta) := (\overline{\alpha} + \theta)u'(\cdot) + v'(\cdot)$ and differentiating twice with respect to θ yields

$$A''(\theta) = 2L^*u''(\cdot) + (\overline{\alpha} + \theta)(L^*)^2u'''(\cdot)\,. \tag{7}$$

Unlike (3), the sign of this derivative cannot be determined even if additional restrictions are imposed on $u(\cdot)$. The reason is that θ appears in the second term of (7), which is positive in some states and negative in others. Hence we cannot obtain clear predictions about how self-employed labour supply responds to a MPS in θ.

What if we utilise a less general definition of increased uncertainty than the MPS?

For example, define a new shock $\tilde{\theta} = \lambda\theta$. Starting from $\lambda = 1$ (where $\theta = \tilde{\theta}$) an increase in λ is a *Sandmo increase in risk*. Substitute $\tilde{\theta}$ for θ with $\lambda = 1$ in the first order condition (6) and partially differentiate with respect to λ to obtain

$$\frac{\partial\{\mathcal{E}[(\overline{\alpha}+\theta)u'(\cdot)]+v'(\cdot)\}}{\partial\lambda} = \mathcal{E}[\overline{\alpha}^2 L^* u''(\cdot)\theta^2] + \mathcal{E}\{\overline{\alpha}u'(\cdot)\theta + [u'(\cdot)+\overline{\alpha}L^* u''(\cdot)]\theta\}\,. \quad (8)$$

The second term on the RHS of (8) can be written as

$$\begin{aligned}&\overline{\alpha}\mathcal{E}(\theta)\{\mathcal{E}[u'(\cdot)] + L^*\mathcal{E}[u''(\cdot)]\} + \mathcal{C}[u'(\cdot),\theta] + \mathcal{C}[\overline{\alpha}L^* u''(\cdot),\theta]\\ &\quad = \mathcal{C}[u'(\cdot),\theta] + \mathcal{C}[\overline{\alpha}L^* u''(\cdot),\theta]\,, \qquad (9)\end{aligned}$$

since $\mathcal{E}(\theta) = 0$. The first term on the RHS of (9) is negative because a larger θ implies a larger y and hence a smaller $u'(\cdot)$. The first term of (8) is also negative by $u''(\cdot) < 0$. But the second term of (9) is positive if $u'''(\cdot) > 0$, as argued earlier, and is only negative if $u'''(\cdot) > 0$. Hence it is only possible to sign the effect of greater uncertainty on labour supply (i.e., as negative) in the 'unrealistic' case of *increasing* absolute risk aversion. But even here, the prediction of lower labour supply in response to greater multiplicative wage uncertainty cannot resolve the labour supply puzzle outlined earlier. And in the more realistic case of non-increasing absolute risk aversion, the above analysis shows that even a Sandmo increase in risk cannot yield any clear-cut results.

The inconclusive nature of this analysis is, not surprisingly, mirrored by previous studies that have analysed similar problems using non-separable utility functions.[4] For example, Block and Heineke (1973) utilised a Slutsky decomposition, and Hartwick (2000) conducted a numerical simulation exercise using a CES utility function. In neither case were unambiguous results discovered, except for a particular special case considered by Hartwick (2000), where stochastic wage outcomes are clustered closely together and individuals' elasticities of substitution are greater than unity. Then a MPS in θ results in *greater* labour supply, as in the case of additive shocks.

What we have shown is that even in the special case of separable utility, it is difficult to obtain any clear prediction about the effects of greater multiplicative wage uncertainty on self-employed labour supply. It seems that the interaction between the stochastic shock and the choice variable is what complicates the multiplicative case relative to the additive one. These theoretical results highlight the value of bringing data to bear on the problem, in order to resolve the matter empirically.

[4] If our utility function was non-separable, then (8) has an extra term $\overline{\alpha}L^*\mathcal{E}(\theta\partial^2 u/\partial y\partial L^*)$, whose sign is state-dependent. No clear-cut result is possible here, even if restrictions are imposed on the third derivative of the utility function.

2.3 An encompassing specification

In empirical applications, it is hard to identify shocks as being either specifically additive or multiplicative. Indeed, it seems likely that both kinds of shock might impinge on earnings simultaneously, for example if output price shocks are positively correlated with shocks to exogenous earnings. And in practice, not only are the sources of income variation typically unobserved, but so also are the shocks themselves. This raises questions about the specification and instrumentation of self-employment wage uncertainty in a labour-supply regression.

Based on the foregoing analysis, a natural candidate for an encompassing income specification is

$$g(wL^* + B; \theta) = \overline{E}_0 + \xi\theta + (\overline{\alpha} + \theta)L^* + B\,,$$

where $\xi > 0$ gives the weight of additive relative to multiplicative shocks. This specification provides a plausible and reasonably general basis for our empirical investigation. We will find it convenient in our investigation to use the standard deviation of past wages as an instrument for uncertainty. This incorporates both additive and multiplicative shocks, as we now show.

Suppose we have sample data for individuals i at time t. If we also allow for the (realistic) possibility that reported self-employed labour earnings are measured with error, ϵ_{it} [where $\mathcal{E}(\epsilon_{it}) = 0$], then the self-employed wage of i at t is computed as follows:

$$\hat{w}_{it} := \frac{\overline{E}_{0it} + \xi\theta_{it} + (\overline{\alpha}_{it} + \theta_{it})L^*_{it} + \epsilon_{it}}{L^*_{it}} = \overline{\alpha}_{it} + \theta_{it} + \frac{\overline{E}_{0it} + \xi\theta_{it} + \epsilon_{it}}{L^*_{it}}\,.$$

Hence although $\mathcal{E}(\hat{w}_{it}) = \overline{\alpha}_{it} + (\overline{E}_{0it}/L^*_{it})$ since $\mathcal{E}(\theta_{it}) = 0$, we have $\mathcal{C}(\hat{w}_{it}, L^*_{it}) < 0$, causing a negative bias if a labour supply equation L^* is estimated as a function of $\hat{w}_{it}$ by least squares.[5] This requires $\hat{w}_{it}$ to be instrumented. We discuss an appropriate choice of instruments in the next section.

3 Empirical specification

Let L_i be the observed number of hours supplied by a self-employed individual i; let w_i be their computed hourly wage (which, as this is a self-employed sample, we will discuss shortly); and let r_i be the standard deviation of i's past log wages – our measure of wage uncertainty, as noted above. Also, let μ_i denote i's unearned income,

[5] This is a (generalised) counterpart of the phenomenon of 'division bias' in the labour supply literature: see, e.g., Rogerson and Rupert (1993).

and let $\mathbf{X}_i$ denote a vector of personal and job-specific characteristics (including a column of ones). Consider the following simple labour supply specification:

$$\ln L_i = \alpha \ln w_i + \beta \ln \mu_i + \gamma r_i + \mathbf{X}_i' \delta + \Upsilon_i + \zeta_i \,, \tag{10}$$

where α, β are γ are scalars, δ is a k-dimensional vector of coefficients, Υ_i are individual fixed effects, and ζ_i is a disturbance assumed to be normally and independently distributed, with mean zero and variance σ_ζ^2.

Several special problems attend the estimation of (10) using sample data on self-employed survey respondents. First, as noted in the previous section, $\ln w_i$ is likely to be endogenous, and so should be instrumented. (A Hausman-Wu test did indeed indicate endogeneity in our empirical application – see Section 5 below.) Second, r_i can only be meaningfully computed for those self-employed individuals who have been observed for a sufficient number of consecutive periods. Consequently there is possible selection bias if these individuals have different characteristics from the rest of the sample. Third, there is also likely to be sampling error in the measurement of r_i, giving rise to a systematic structure in the variance-covariance matrix of the error term. We now take each of these points in turn.

There are several ways of instrumenting $\ln w_i$. One is to propose a set of instruments drawn from the same period in which the wage is observed. The problem with this approach, which is necessary if only cross-section data are available, is that it requires more or less *ad hoc* exclusion restrictions on the set of instruments and the set of regressors, $\mathbf{X}$. In contrast, the availability of panel data enables the researcher to side step this problem by using previous period log wages as instruments. These can be expected to satisfy the usual criteria for valid instruments, namely correlation with $\ln w_i$ but independence from ζ_i. We will adopt this approach below and test the quality of these instruments.

The second consideration is the sample frame. Because r_i is measured as the standard deviation of wages observed in each of n_i previous periods – where n_i can vary from individual to individual – at least two previous observations on each i's earnings are needed to compute r_i. Hence no individual with less than three consecutive wage observations (one current and two previous) can be included in the sample.[6] This raises the possibility of having an unrepresentative, self-selected sample. We will check this in two ways below, both by comparing the sample characteristics of included and excluded respondents, and by incorporating a selectivity correction into the empirical model.

[6]The stipulation that these observations are *consecutive* is an important one, because the concept of specifically *self-employment wage uncertainty* for an individual who mixes self-employment with employment spells does not make obvious sense. We leave the more complicated problem of wage uncertainty defined over mixed spells for future work.

The problem of heteroscedasticity arises because sample values of wage uncertainty, denoted by $\tilde{r}_i$, differ systematically from population values denoted by r_i. Assume that the difference is given by a mean-zero normally distributed error, ν_i. This error recognises that we possess only a (stochastic) sample estimate of non-stochastic r_i, rather than the population value based on the entire (but unobserved) set of wages. That is, $\tilde{r}_i = r_i + \nu_i$, with $\mathcal{E}(\tilde{r}_i) = r_i$, i.e., $\tilde{r}_i$ is an unbiased estimator of r_i. Then the empirical counterpart to (10) is

$$\ln L_i = \alpha \ln w_i + \beta \ln \mu_i + \gamma r_i + \mathbf{X}_i' \delta + \Upsilon_i + (\zeta_i + \gamma \nu_i) \,. \tag{11}$$

Define $\Xi_i := \zeta_i + \gamma \nu_i$. Clearly $\mathcal{E}(\Xi_i) = 0$ and $\mathcal{V}(\Xi_i) = \sigma_\zeta^2 + \gamma^2 \mathcal{V}(\nu_i)$. It can be shown using standard results from intermediate statistics (e.g., Wetherill, 1981, Sec. 1.5) that

$$\mathcal{V}(\nu_i) = \mathcal{V}(\tilde{r}_i) = 2r_i^4/(n_i - 1) \,.$$

Hence it follows that the variance of the disturbance term of (11) can be consistently estimated as

$$\varsigma_i := \widehat{\mathcal{V}(\Xi_i)} = \hat{\sigma}_\zeta^2 + \frac{2\hat{\gamma}^2 \hat{\tilde{r}}_i^4}{n_i - 1} \,. \tag{12}$$

This expression implies that the longer an individual's wages have been observed (i.e., the greater is n_i), the more precise the estimated variance of those wages, as expected.

The above arguments suggest that, after instrumenting the log wage to get $\ln w_i^I$, a two-stage feasible GLS estimator can be used to estimate the key parameters of interest, namely $\Theta := (\alpha, \beta, \gamma, \delta')'$. The GLS estimator is

$$\hat{\Theta} = \left[\sum_{i=1}^{n} \left(\frac{1}{\varsigma_i^2} \right) \mathbf{z}_i \mathbf{z}_i' \right]^{-1} \left[\sum_{i=1}^{n} \left(\frac{1}{\varsigma_i^2} \right) \mathbf{z}_i \ln L_i \right] \,, \tag{13}$$

where $\mathbf{z}_i$ are the columns of $\mathbf{Z}_i := [\ln w_i^I, \ln \mu_i, r_i, \mathbf{X}_i]$. The first step is to estimate the parameters of (11), and the second is to use the results to compute the ς_i and so estimate Θ from (13). It is well known that this is a consistent estimator, and that iterating the above procedure does not improve on it and is generally not recommended in practice (Greene, 2002, Sec. 11.6).

4 Data sample

The data for this study are collected from the Panel Study of Income Dynamics (PSID), using a panel of 26 waves of cleaned observations, from 1968 to 1993. Our sample comprises self-employed male working-age household heads. The self-employed were defined as individuals who were working at the time of the interview; who were self-employed in their main job; and who were not also working for someone else contemporaneously. The sample was restricted to self-employed individuals who earned positive incomes and worked positive numbers of hours (but no greater than 4000 annually), who had at least three consecutive years of income data, and single continuous spells in self-employment (the latter to avoid contaminating estimates of the effects of risk with occupational switching). This left a sample of 1456 observations on 286 self-employed individuals.

The average annual work hours in the sample was 2,555. This is comparable with those reported in other studies, e.g., Carrington *et al* (1996), whose sample average over a similar period was 2,396. Hourly wages were computed by dividing annual earned income by annual hours worked. Both earned and unearned incomes were deflated to 1975 dollars using the CPI. The average hourly wage rate was about \$6.46, which is similar to the rate reported by Hamilton (2000), calculated from 1980s Survey of Income and Program Participation data.[7] The variable $\ln w_i$ was defined as the natural logarithm of one plus the hourly wage. The variable $\ln \mu_i$ was defined as the natural logarithm of one plus the respondent's investment income plus wife's income, if any. The variable r_i was calculated as the natural logarithm of the standard deviation of wages in the immediately preceding contiguous years.

Data were also collected for several explanatory variables, $\mathbf{X}$. These included the number of years in the present job ('tenure'), a dummy for (self-assessed) disability status, age and its square, and the number of children in the household. Summary statistics are collected in Table 1. The average age of sample respondents is a little over 46 years, which partly accounts for the relatively small average number of children in the household, and the relatively long average tenure in the job, of just under 14 years.

As a preliminary check for possible selection bias, we re-computed Table 1 for *all* self-employed individuals with positive incomes and work hours. The summary statistics are reported in Table A1, in the Appendix. These statistics reveal that the only major difference compared with Table 1 is the average tenure being over 5 years

[7] In contrast, Carrington *et al* (1996) used a different sample frame, which also included respondents with employment income. This might partly account for the higher average wage rates reported in that study compared to ours. We also tried adjusting the data for topcoding, by multiplying maximum labour incomes by 1.33 as in Devereux (2003), but this made an imperceptible difference to the results.

Table 1: Summary statistics

Variable names	Description	Mean	Std.Dev.	Minimum	Maximum
LHOURS	Log of annual work hours	7.78	0.43	2.48	8.29
WAGE	Log of real wage	1.60	0.65	0.01	7.55
RISK	Std.Dev. of log past wages	0.21	0.43	0.00	6.25
UNINC	Log of real unearned income	6.57	3.81	0.00	11.53
TENURE	No. years in the current job	13.85	9.77	0.08	50.00
DISAB	Dummy variable = 1 if disabled	0.10	0.30	0.00	1.00
AGE	Age	46.60	10.80	21.00	65.00
AGE^2	*AGE* squared/100	22.89	9.92	4.41	42.25
NCHILD	No. of children in the household	1.18	1.28	0.00	7.00

Source: PSID, 1968–93, authors' calculations. Maximum number of observations: 1456. Minimum number of observations: 1379.

longer in the latter. This is of course just as we would expect, reflecting our need to focus on longer self-employment spells. In all other respects, the characteristics of the narrower sample are fairly similar to those of the broader sample. Hence we do not have strong *a priori* grounds for believing that selection bias will be a problem.

5 Results

We first tested whether the computed wage is endogenous. A Hausman-Wu statistic of $\chi^2(1) = 4.20$ indicated that it is. As explained in Section 3, we instrumented the wage using the previous two years' wages, together with individual fixed effects. An F statistic of $F(2, 1156) = 11.727$ confirms the quality of these instruments.

We first estimated a model of labour supply without the risk variable, termed model (1). We then re-estimated the model including the risk variable, to gauge directly the effect of omitting it. This is termed model (2). The results are collected in Table 2.

The results for model (1) suggest that the wage rate is a key explanatory variable, with a significant negative effect. This is indicative of a backward bending labour supply schedule, and mirrors the findings of other authors for self-employed workers.[8]

[8]E.g., Wales (1973) found evidence of backward-bending labour supply behaviour by American

Table 2: Estimates of the self-employed labour supply model

Variable	**Model (1)**	**Model (2)**
WAGE, $\ln w^I$	-0.234*	-0.113
	(0.102)	(0.107)
RISK		0.085**
		(0.023)
UNINC	0.004	0.002
	(0.003)	(0.003)
TENURE	-0.001	-0.001
	(0.001)	(0.001)
DISAB	0.032	0.014
	(0.030)	(0.030)
AGE	0.117	0.088
	(0.117)	(0.116)
AGE^2	-0.025*	-0.022
	(0.012)	(0.012)
NCHILD	0.009	0.010
	(0.010)	(0.010)
No. Obs.	1379	1379
R^2	0.80	0.80
$F(k, n-k)$	15.35**	59.00**

Source: Dependent variable is the natural logarithm of the annual hours of work. Results include individual fixed effects (not reported). The estimator is two-stage feasible GLS, with an instrumented wage. Instruments are two period lagged wages. Standard errors are in parentheses. * p-value less than 0.05; $^{**}p$-value less than 0.01.

The only other significant variable is the square of age, which also has a significant negative effect. This implies that self-employed individuals adjust their work hours downwards as they age, with older workers adjusting them more than their younger counterparts do. Evidence of this kind of behaviour is quite well established in work on partial retirement by the self-employed (see, e.g., Fuchs, 1982; Honig and Hanoch, 1985).

The results for model (2) question the robustness of these results. When the risk variable is included, the significant effects from the wage rate disappear. The effect from squared age is a little smaller and no longer significant at a 5 per cent significance level (though it is significant at the 10 per cent level). In contrast, risk exerts a significant positive effect on work hours. It is also the only significant determinant of self-employed work hours. In line with our earlier theoretical analysis, this finding is consistent with the self-employed self-insuring to compensate for greater risk. While it is unclear whether this reflects a dominant role for additive relative to multiplicative risk, it does have two important implications. First, risk appears to be the key to understanding self-employed labour supply. And second, the positive effect on labour supply can in principle explain the self-employed labour supply puzzle outlined in the Introduction (see Section 2.1).

One might ask why the inclusion of risk in the labour-supply specification renders the wage rate insignificant. One answer might be along standard econometric lines: that the exclusion of a relevant variable biases the coefficients (and standard errors) of all other variables in the regression (Greene, 2002). Thus, the apparent finding (in model (1) and in the previous self-employed labour supply studies cited in footnote 8) of a significant negative effect from the wage rate might be spurious. While possibly appealing, this answer does however beg the question as to why the estimated coefficient on the instrumented wage becomes closer to zero, rather than diverging from zero, when r is included. We cannot readily suggest any reasons based on the theoretical analysis as to why this might be so. But we can probably rule out an explanation based on multicollinearity between the instrumented wage and the standard deviation of past wages. Collinearity does not bias regression coefficients so cannot explain the differences in the estimated coefficient on the instrumented wage in columns (1) and (2) of Table 2. While collinearity does increase estimated standard errors, Table 2 offers little evidence that this has occurred either; and the partial correlation coefficient between r and $\ln w^I$ is only 0.113 in any case.

We went on to conduct two robustness checks on the empirical results. In the first, we measured r as the variance (rather than the standard deviation) of previous

self-employed proprietors using NBER-Thorndike data. Camerer *et al* (1997) also reported negative effects from wage rates on the labour supply of self-employed New York taxi-drivers. They took this as evidence that taxicab drivers 'take one day at a time'.

log wages. The results were virtually unchanged from those of column (2), with a γ estimate of 0.124 (standard error = 0.053). Hence we conclude that the results are robust to a different functional form of the risk variable.

A second issue that needs to be addressed is that of sample selection. In this instance, as we have already mentioned, the possible sample selection is that those self employed individuals with the 3 or more years of observed self-employment incomes necessary to compute our risk measure are a non random selection from the wider sample of self-employed. As we have already noted, a comparison of Table 1 with Table A.1 indicates that the two samples are similar. However, to check this further we added to the estimated hours equation a Heckman correction term evaluated at the estimates obtained from a fixed effects probit with dependent variable equalling one for self employed individuals with 3 or more years observed self employed income, and zero with 2 or less. This probit contains full year dummies and is reported as Table A.2 in the Appendix. The coefficient on the included selection term was insignificant with value 0.016 and standard error of 0.043. We conclude that while the self-employed themselves may be a self-selected sample, those remaining self-employed for 3 or more consecutive periods does not involve a further selection on characteristics, and so we can base inference regarding the determinants of self employed labour supply with some confidence on this sample.

6 Conclusion

This paper presented a new approach to thinking about the labour supply behaviour of self-employed individuals. We developed a theoretical model that explores the different ways that wage uncertainty might affect self-employed labour supply. An empirical investigation based on PSID data suggested that risk, proxied by the standard deviation of past log wages, is the sole determinant of self-employed labour supply, and its inclusion in a self-employed labour supply regression model rather strikingly rendered the previously significant wage insignificant. In addition, these results might also supply an answer to the long-standing puzzle of why the self-employed work such long hours for relatively low wages. Of course, these findings have only been derived from one data set, and further work using other data is needed. But they do suggest that further work on the topic of labour supply under wage uncertainty generally, and of self-employed labour supply in particular, might usefully clarify and extend our knowledge about these important topics.

References

Aronson, R. L. (1991) *Self Employment: A Labour Market Perspective*, ILR Press, Ithaca, New York.

Arrow, K. J. (1971) *Essays in the Theory of Risk Bearing*, Markham Press, Chicago.

Block, M. K. and J. M. Heineke (1973) The allocation of effort under uncertainty: the case of risk-averse behaviour, *Journal of Political Economy*, **81**, pp. 376–85.

Blundell, R. and T. MaCurdy (1999) Labour supply: a review of alternative approaches, in *Handbook of Labour Economics, Volume 3A* (eds. O Ashenfelter and D. Card), Elsevier, Amsterdam, pp. 1559-1695.

Boulier, B. L. (1979) Supply decisions of self-employed professionals: the case of dentists, *Southern Economic Journal*, **45**, pp. 892–902.

Camerer, C., Babcock, L., Loewenstein, G. and R. Thaler (1997) Labour supply of New York City cabdrivers: one day at a time, *Quarterly Journal of Economics*, **112**, pp. 407–41.

Carrington, W. J., McCue, K., and B. Pierce (1996) The role of employer/employee interactions in labour market cycles: evidence from the self-employed, *Journal of Labour Economics*, **14**, pp. 571–602.

Devereux, P. J. (2003) Changes in male labour supply and wages, *Industrial and Labor Relations Review*, **56**, pp. 409–28.

Evans, D. S. and L. S. Leighton (1989) Some empirical aspects of entrepreneurship, *American Economic Review*, **79**, pp. 519–35.

Fuchs, V. R. (1982) Self-employment and labour-force participation of older males, *Journal of Human Resources*, **17**, pp. 339–57.

Greene, W. H. (2003) *Econometric Analysis*, Fifth Edition, Prentice-Hall, New Jersey.

Hamilton, B. H. (2000) Does entrepreneurship pay? An empirical analysis of the returns to self-employment, *Journal of Political Economy*, **108**, pp. 604–31.

Hartwick, J. M. (2000) Labour supply under wage uncertainty, *Economics Letters*, **68**, pp. 319–25.

Honig, M. and G. Hanoch (1985) Partial retirement as a separate mode of retirement behaviour, *Journal of Human Resources*, **20**, pp. 21–46.

Kihlstrom, R. E. and J. J. Laffont (1983) Taxation and risk taking in general equilibrium models with free entry, *Journal of Public Economics*, **21**, pp. 159–81.

Kraus, A. and R. Litzenberger (1976) Skewness preference and the valuation of risk assets, *Journal of Finance*, **31**, pp. 1085–1100.

Parker, S. C. (1997) The distribution of self-employment income in the United Kingdom, 1976–1991, *Economic Journal*, **107**, pp. 455–66.

Parker, S. C. (2004) *The Economics of Self-employment and Entrepreneurship*, Cambridge, Cambridge University Press.

Phillips, B. D. and B. A. Kirchhoff (1989) Formation, growth and survival: small firm dynamics in the US economy, *Small Business Economics*, **1**, pp. 65–74.

Rogerson, R. and P. Rupert (1993) On testing the intertemporal substitution theory of labour supply, *Journal of Economic Dynamics and Control*, **17**, pp. 37–50.

Schaffner, J. A. (1993) Rising incomes and the shift from self-employment to firm-based production, *Economics Letters*, **41**, pp. 435–40.

Scheffler, R. M. and L. F. Rossiter (1983) Compensation schemes and the labour supply of dentists, *Quarterly Review of Economics and Business*, **23**, pp. 29–43.

Taylor, M. P. (1999) Survival of the fittest? An analysis of self-employment duration in Britain, *Economic Journal*, **109**, pp. C140–C155.

Thornton, J. (1998) The labour supply behaviour of self-employed solo practice physicians, *Applied Economics*, **30**, pp. 85–94.

Wales, T. J. (1973) Estimation of a labour supply curve for self-employed business proprietors, *International Economic Review*, **14**, pp. 69–80.

Wetherill, G. B. (1981) *Intermediate Statistical Methods*, Chapman and Hall, London.

Appendix

Table A1. Summary statistics for all self-employed workers

Variable names	**Mean**	**Std.Dev.**	**Minimum**	**Maximum**
LHOURS	7.63	1.80	1.00	8.29
WAGE	1.60	0.77	0.01	7.56
UNINC	5.95	3.74	0.00	11.53
TENURE	8.42	8.83	0.08	60.67
DISAB	0.12	0.32	0.00	1.00
AGE	41.89	11.15	18.00	65.00
AGE^2	18.79	9.68	3.24	42.25
NCHILD	1.26	1.37	0.00	10.00

Source: PSID, 1968–93, authors' calculations. Maximum number of observations: 10319. Minimum number of observations: 9531.

Table A2. Fixed effect probit model with year dummies

Variable	Probit
YRSED	-0.006
	(0.015)
TENURE	0.037**
	(0.003)
DISAB	0.128
	(0.083)
AGE	5.010**
	(0.358)
AGE^2	-0.386**
	(0.029)
EDTEN	0.825**
	(0.189)
No. Obs.	9531
$-LL$	3401.69

Source: Dependent variable is one if the individual was self-employed for three or more years consecutively, and zero if self-employed for two consecutive years or less. YRSED is years of education; EDTEN is the interaction of YRSED and TENURE. Results include individual fixed effects and time dummies (not reported). Standard errors are in parentheses. For asterisks, see Table 1.