UNITED KINGDOM · CHINA · MALAYSIA

Discussion Papers in Economics

Discussion Paper
No. 13/03

Success Taxes and Entrepreneurship: A Regression Discontinuity Approach

Richard Kneller and Danny McGowan

October 2013

Success Taxes and Entrepreneurship: A Regression Discontinuity Approach

Richard Kneller† and Danny McGowan‡*

Abstract

We study the causal effect of upside progressivity in the personal income tax schedule, known as success taxes, on entrepreneurship in the UK by exploiting sharp discontinuities in the convexity of taxation that arise at government-specified thresholds in the income distribution. Using a regression discontinuity design we find no evidence that assignment to greater upside convexity around the top tax bracket affects transitions into self-employment. This result is robust to examining the question over longer time horizons, using other threshold points, and cannot be explained by offsetting jumps in other covariates or strategic behaviour.

Keywords: Self-employment, taxation, convexity, entrepreneurship

JEL Codes: H2, H3, D2, L2, J2

† Richard Kneller, University of Nottingham, School of Economics, University Park, Nottingham, NG7 2RD, United Kingdom. Email: richard.kneller@nottingham.ac.uk; tel: +44 (0)115 951 4734; fax: +44 (0)115 951 4159.

‡ Danny McGowan, Bangor University, Business School, Hen Goleg, College Road, Bangor, LL57 2DG. Email: d.mcgowan@bangor.ac.uk; tel: +44 (0)1248 38 3948.

* We are grateful for comments and suggestions by Marius Brülhart, Norman Gemmell and Ismael Sanz, and helpful discussions with conference participants at the 2012 Royal Economic Society Annual Conference.

1. Introduction

The important employment, growth and innovation effects associated with entrepreneurship have raised questions within the public policy literature of whether governments can, or should, affect the amount or the type of entrepreneurship that takes place within an economy.[1] One feature of this debate has been the response of entrepreneurs to tax policy. Traditionally, the empirical evidence used to support these arguments has focused on the effects of the level of, or changes to, marginal and average tax rates (see for example Carroll et al. 1998, Bruce, 2000; Bruce and Deskins, 2012; Djankov et al. 2010). More recently, recognising the risk entailed in entrepreneurship, the literature has begun to explore the role played by non-linearities in the tax system.[2]

The discrete nature of starting a new business, combined with its associated risks, indicates that tax policy can affect the expected returns to this investment, and therefore non-linear tax systems play a potentially important role in entrepreneurial decisions. With progressive marginal tax rates and imperfect loss-offset provisions Gentry and Hubbard (2000, 2005) show that successful entrepreneurs face a higher average tax rate than unsuccessful ones, which reduces the incentive to become self-employed. A risk neutral investor will therefore require a higher pre-tax expected rate of return on risky projects, discouraging risk-taking. Using micro data for the US they find that progressivity, measured by the difference in the expected marginal tax rate (MTR) paid in successful and unsuccessful states of entrepreneurship is negatively correlated with entry into (risky) business ownership.[3] Wen and Gordon (2013) construct an alternative measure of tax convexity based on the expected value of the tax liability of an entrepreneur facing a distribution of possible returns compared to their tax liability at their predicted income. The structural approach they employ indicates that progressive income tax systems discourage entrepreneurship in Canada.[4]

Whilst Gentry and Hubbard (2000, 2004, 2005) and Wen and Gordon (2013) find a similar significant effect from tax convexity on self-employment decisions, both note the potential for an asymmetry in the effect of tax convexity in the Gentry and Hubbard (2000, 2004, 2005)

[1] See for example Bartelsman et al. (2005), Baumol (2002), OECD (2001) and Disney et al. (2003).
[2] See for example Gentry and Hubbard (2000, 2004, 2005), Cullen and Gordon (2007), Da Rin et al. (2011), Bacher and Brülhart (2013), Wen and Gordon (2013), and Ferede (2013).
[3] The remaining studies test for the effects of convexity in corporate taxation on firm-level outcomes. Da Rin et al. (2011) and Bacher and Brülhart (2013) study the effects of progressivity of the corporate tax schedule on forms of new firm creation for 17 European countries and Swiss provinces respectively. Both find a significant positive effect, although Bacher and Brülhart (2013) conclude that progressivity is of secondary importance to the level of taxation. Cullen and Gordon (2007) argue that successful entrepreneurs can avoid high personal income tax rates by incorporating and show using US data that the greater the difference between personal and corporate tax rates, the greater is the incentive to undertake riskier investments. The exception is Ferede (2013) who includes a measure of progressive personal income taxes as a control variable using Canadian-province level data on rates of self-employment. He finds a significant negative effect at high income levels (167% of average wages) but not at lower income levels (100% of average wages).
[4] There is also a larger literature on of the effects of tax convexity on other firm-level decisions such as hedging using financial derivatives (see Graham and Smith, 1999 and Graham and Rogers, 2002 as examples), investment and exit (Sarkar and Goukasian, 2008), and corporate finance (Sarkar, 2008).

measure. Wen and Gordon (2013) argue that within a measure based on the spread of taxes it is clear why high-tax rates in the event of success may discourage risk-taking, but it is less clear why low tax rates in the event of failure discourage it.[5] Consistent with this, Gentry and Hubbard (2000, 2004, 2005) find empirical evidence that the effects of tax convexity are stronger when the convexity faced by an individual is on the upside rather than downside.[6] This result helps explain their labelling of the effects of progressive tax schedules on self-employment as 'success taxes'.

The objective of this paper is to assess the extent to which success taxes affect the decision to become an entrepreneur in the UK (measured through transitions from employment into self-employment).[7] For this task we exploit the large changes in success taxes that occur around an exogenous government-specified income threshold that assigns employees to the top bracket (a marginal tax rate of 41%) or a lower bracket (with a marginal tax rate of 23%). As we discuss in more detail in the next section of the paper, at this threshold the Gentry and Hubbard (2000, 2004, 2005) construct of upside convexity turns from positive to zero. The changes in downside convexity that occurs exactly offset this, such that total convexity does not change. It follows that if 'success taxes' are important for risky-entrepreneurship decisions, then individuals with a current wage income just below the tax threshold face positive upside convexity should be less likely to enter self-employment compared to those with a current wage income just above the tax threshold, who face no upside convexity.[8] Importantly, this discontinuity in upside and downside convexity occurs independently of the assumed distribution of expected returns from self-employment used in the construction of the Gentry and Hubbard (2000, 2005) measure: the same change occurs for all individuals. It follows that this feature of tax convexity can be exploited as a means to identify the causal effects of success taxes on entrepreneurship for the UK, or indded other countries where there are jumps in the marginal income tax rate. Accordingly, we adopt for this task a regression discontinuity (RD) design (Lee, 2008; Lee and Lemieux, 2010).

To answer this question we use individual-level microdata from the UK British Household Panel Survey between 2000 and 2007. Before employing parametric and non-parametric RD methods for our main tests, we carefully demonstrate observational equivalence in terms of observable characteristics between individuals in the treatment (low tax bracket) and control (higher tax bracket) groups. Importantly, we also show that individuals are unable to influence

[5] The measure used by Wen and Gordon (2013) captures the surtax on risky income and is therefore not affected by such arguments.

[6] Upside convexity occurs when expected success as an entrepreneur increases the individual's marginal tax rate relative to the baseline case (working for someone else), while downside convexity occurs because being an unsuccessful entrepreneur lowers it.

[7] Self-employment is commonly used as a measure of entrepreneurship within the literature, albeit one that is recognised does not fully capture all forms of entrepreneurship. Data limitations prevent us from using alternative measures.

[8] The Wen and Gordon (2013) construct of tax convexity does not display the same discontinuity at the top tax threshold (see for example Figure 1 in their paper).

treatment status and that clustering on one side of the threshold, due to strategic tax avoidance behavior, is absent from our data.[9]

We find from our analysis *no* evidence in support of the hypothesis that upside convexity of the income tax schedule affects decisions to enter self-employment. Facing greater upside convexity in the personal income tax schedule has an insignificant effect on the probability that an individual will become self-employed the next year, or indeed at any point over the next 5-7 years. We show that these findings are not a consequence of using the top tax threshold: the same result is found when using an RD design around other thresholds in the income tax schedule.

We explore two alternative explanations for the difference in our findings compared to Gentry and Hubbard (2000, 2004, 2005).[10] Firstly, alongside the difference in the success taxes faced by an individual at the top tax threshold, there is obviously also a large difference in the marginal income tax rate itself. A number of papers within the literature, including those by Long (1982), Blau (1987) Schuetze (2000), Gentry and Hubbard (2000), have shown using a combination of macro and micro data that self-employment is increasing in the marginal tax rate, although Fölster (2002), Parker (2003) Bruce and Mohsin (2006) have found no effect, or a negative effect. The positive effects of marginal income tax rates are usually interpreted as supportive of the view that self-employment offers increased opportunities to evade taxation through both legal and illegal means. Self-employment income can be under-reported in a way that is more difficult for employees. Or, the self-employed can avoid taxation through legal opportunities to deduct business related consumption from taxable income. The presence of such effects from higher marginal tax rates would be expected to move in the same direction as the effects from facing greater convexity in tax rates. To put this another way, had we found a negative effect of paying higher marginal tax rates, tax evasion/avoidance might provide an alternative explanation for our results. As we do not, our results also provide evidence against the commonly held view that employees are motivated to become self-employed in order to evade or avoid higher marginal income tax rates.

The second explanation builds on Bacher and Brülhart (2013) who argue that the results for measures of tax progressivity, in their case convexity of corporate taxes, are affected by the omission of other relevant tax variables. Given the discrete nature of the investment of starting a new business, alongside tax convexity, conceptually the average tax rate should be a key factor behind the decision to become self-employed. We test the sensitivity of our results to the inclusion of four measures of the average tax rate and the Wen and Gordon (2013) measure of

[9] Saez (2010) also finds that, with the exception of the personal allowance threshold, tax payers in the US do not bunch at kink points in the income distribution.

[10] We also attempted to replicate the results of Gentry and Hubbard (2000, 2004, 2005) using a probit analysis across all individuals. Using a similar method of calculating progressivity of marginal tax rates to Gentry and Hubbard (2000, 2004, 2005) we found an insignificant effect from total tax convexity and no evidence that the effects are asymmetric. These results are available from the authors on request.

convexity. The average tax rate measures include the average tax rate paid in employment,[11] a measure of the difference in the average tax rate paid in employment and self-employment at an individual's current income level,[12] the probability weighted expected average tax rate in self-employment,[13] and finally, developing the idea that there may be asymmetries in the effects of taxation, we also include a variant of this measure that uses just the expected average tax rate if they were to become a successful entrepreneur.

We find some role for omitted tax terms, but not alternative measures of convexity.[14] We find that the difference in the average tax rate paid in employment and self-employment has some explanatory power in at least some of our regressions, and has the expected relationship with transitions into self-employment. The greater the difference in the average tax rate paid in employment versus self-employment, the greater is the probability to move into self-employment. We find no significant effects from the other tax rates that we study.

The rest of the paper is structured as follows. In the next section of the paper we outline in more detail the expected relationship between progressivity of the tax schedule and entrepreneurship and estimates of the convexity in the UK tax schedule. We use this to motivate the empirical approach that we adopt in the paper. Section 3 provides details on the data, while Section 4 provides detail on the empirical methodology. The main results and the various robustness tests and extensions are provided in Section 5. Conclusions are drawn in Section 6.

2. Convexity and the Income Tax Schedule

To make clear the reasons behind the empirical methodology that we employ, we provide some detail on convexity in the UK personal tax system, and upside and downside convexity. The starting point for this discussion is the Gentry and Hubbard (2004, 2005) model of entrepreneurial selection. In their model entrepreneurship is a discrete choice affected by average tax rates as potential entrepreneurs compare the after-tax returns of remaining in employment versus working themselves and paying tax as a self-employee. Encompassing the feature that income tax schedules are typically progressive in nature they show how this non-linearity acts to tax successful projects at a higher rate than unsuccessful ones. As risky projects have a higher variance of returns, progressivity of the tax schedule acts to discourage risk taking, even by risk-neutral individuals. In addition to the effects for total convexity, evidence of an asymmetric effect from the progressivity of taxation on entrepreneurship is important for

[11] This draws on evidence from Robson and Wren (1999), Schuetze (2000), Bruce (2000) and Cullen and Gordon (2002).
[12] This is motivated by the work of Bruce (2000) and Schuetze (2000) and Fairlie and Mayer (1999).
[13] This follows from Bacher and Brülhart (2013).
[14] It is worth pointing out that this does not imply that tax convexity or the average tax rate do not matter for the decision to become self-employed. As there are no discontinuities in the values for these variables around the higher tax rate threshold our regressions have weak explanatory power for these variables. We conclude only that they are not a confounding factor in our regressions.

our understanding of the design of tax systems, as they imply that there are additional points in the tax schedule at which behavioural changes will occur, which may have important welfare effects.

The empirical counterpart of their measure of convexity compares the weighted average of the marginal tax rates in various successful and unsuccessful states of entrepreneurship with that of remaining in employment.[15] To generate this measure for each taxpayer (households in the US data) they consider successful and unsuccessful states as a multiple of current income levels and assign probability weights to these. The incomes of successful entrepreneurs are assumed to increase current labour income by 25, 50, 100 and 200%, where probability weights of 0.5, 0.3, 0.15 and 0.05 are assigned to each of those respectively. A similar method is applied to unsuccessful states, but where current income falls by 10, 25, 50 or 75%. These are given probabilities of 0.5, 0.3, 0.15 and 0.05 respectively. As a reminder, upside convexity is measured as the difference between the marginal tax rate when successful compared to remaining in employment and downside convexity as the difference in the marginal tax rate in self-employment and employment when unsuccessful. The overall measure of convexity is the sum of these two values.

The corresponding marginal income tax rate, and therefore convexity, will of course vary with the assumptions regarding the distribution of expected returns were an individual to become self-employed. Based on these estimates, Gentry and Hubbard (2004) discuss in detail how the estimated convexity faced by an individual can vary between positive and negative values, and according to their personal circumstances (the tax allowances they can access). They also describe how they can vary considerably for the same individual across time as various tax rates and thresholds are altered. Using data on UK tax rates and thresholds, and the same assumptions about the distribution of expected returns as Gentry and Hubbard (2000, 2004, 2005), we show that this is similar in the UK.

We begin with some discussion of personal income taxation in the UK. Income taxes for the employed and self-employed in the UK are a combination of personal income taxes and social security contributions (national insurance). The personal income tax operates through a system of allowances and bands of income. Over the period in which we study there are four tax bands in most years. Each individual has a personal allowance which is taxed at a zero marginal rate, a lower rate of 10%, a basic rate of 22% and a top marginal tax rate of 40%. As an example, in

[15] In this paper we follow Gentry and Hubbard (2000, 2004, 2005) in using progression of marginal income tax rates. Since Musgrave and Tun Thin (1948) it has been known that other measures of progressivity of the tax schedule are also possible. Of the most commonly used alternative measures of progressivity, a discontinuity would be present for the residual income progression measure due to Musgrave and Musgrave (1989). This is calculated as the ratio of 1- marginal tax rate over 1- average tax rate. A discontinuity would not be present for progressivity of average tax rates however. Using the assumed distribution of expected returns from self-employment as for marginal income taxes we have calculated estimates of convexity in average tax rates. These do not display the same discontinuity around the top tax threshold as the measure of convexity in marginal tax rates and have no predictive power when added to the regression. These results are available from the authors on request. We consider a more recent alternative measure of progressivity due to Wen and Gordon (2013) in section 5.3.

2007 employees pay the top marginal income tax rate when taxable income reached £34,600 (when earnings reach £39,825 when combined with the personal allowance).

National Insurance (NI) contributions have a different set of income bands and marginal tax rates. Employees pay NI contributions at a rate of 11% (figure for 2007) between a primary threshold of £100 (weekly gross earnings) and an upper threshold of £670 (weekly gross earnings). Below the primary threshold the marginal tax rate is zero and above the upper threshold the rate drops to 1%. In 2007 the upper threshold was £34,840. The self-employed pay a somewhat different set of NI contributions. They pay a flat rate of £2.20 per week (2007 figures) when their income exceeds the smaller earnings exception threshold and then 8% on any profits between the lower and upper threshold. This rate also drops to 1% above the upper threshold. The usual explanation given for the lower rates of national insurance contributions by the self-employed is that the latter have less benefit entitlement. In the UK when a person has made sufficient NI contributions they may claim unemployment benefit (known as higher rate job seekers allowance), incapacity benefit, old-age pension etc. The self-employed have lower benefit entitlements, especially with regard to unemployment benefit.

The difference in the threshold income levels for personal income taxation and NI contributions leads to a fairly complicated set of marginal tax rates for employees and the self-employed set out in Figure 1. These jumps occur because the personal income tax thresholds are located at different income levels to those for NI. For example, there is a fall in the combined marginal tax rate for incomes between £34,840 and £39,825 for both the self-employed and employees. For incomes just below £34,840 the marginal income tax is 33% for employees and 30% for the self-employed. This falls to 23% for both groups when income is between £34,840 and £39,825, before jumping again to 41% for income beyond the top threshold. Table 1 shows the corresponding combined tax rates and thresholds for employees and the self-employed for the years from 2000 to 2007.

[INSERT TABLE 1] [INSERT FIGURE 1]

In Figure 2 we plot the estimated values for total convexity alongside upside and downside convexity using the same values for the distribution of expected returns to self-employment as Gentry and Hubbard (2000, 2004, 2005) but using the UK marginal tax rates reported in Table 1.

As is clear from Figure 2 even for a single year there is considerable variation in both upside and downside convexity according to the income of the individual, with large positive peaks in convexity at low income levels, around the personal tax allowance, and at around £30-40,000, when NI contributions drop to 1%, and the personal income tax rate jumps from 22% to 40% (so that the combined rate jumps from 23% to 41%). At no point in the income schedule do we calculate there to be negative values for total convexity, but there are a number of values of wage income where this occurs for downside convexity. In Figure 3 we show estimates of convexity applying the UK tax code for 2001, 2004 and 2007. As in the US case it is clearly possible for individuals to face markedly different estimates of convexity across years even if their income did not alter.

In both Figure 2 and Figure 3 it is evident that there are discrete jumps in the values of total convexity at many points in the tax schedule. Where these occur is dependent upon the location of the tax thresholds and marginal tax rates but also the assumed distribution of expected returns to self-employment (the construction of this figure assumes that income is expected to fall or rise by the fixed multiples of current income used by Gentry and Hubbard and discussed already above). Changes to total convexity occur because of changes to upside and downside convexity. Of interest to the empirical design we adopt, upside and downside convexity change on a more frequent basis than total convexity and in particular at points close to thresholds in the combined income tax schedule.

[INSERT FIGURE 2] [INSERT FIGURE 3]

As described in the introduction, our focus is on the values of convexity around the top tax threshold. In Figure 4 we plot upside and downside convexity for income between £35,000 and £45,000 using 2007 tax rates, which is one of the broadest income windows we use in our regressions. As a reminder, the top income tax threshold starts when income reaches £39,825 in that year. Upside convexity has a value of 18 up until the top tax threshold and then drops to zero, whereas downside convexity has the value of -7, then -3.5 close to the threshold and 14.5 above the threshold. Total convexity is therefore 11, before increasing to 14.5 once wage income is close to and above the threshold.[16] For incomes close to the threshold, the total convexity faced by an individual is constant but whether this is upside or downside convexity is not. Those below the top tax threshold have greater upside convexity, and therefore are subject to greater success taxes, compared to those above the top tax threshold, who face no upside convexity. This is the crucial insight that forms the heart of the identification strategy that we adopt in the paper.

[INSERT FIGURE 4]

A key question here must be how general is this point about the values of upside and downside convexity at the top threshold of marginal personal income tax rates. It can be considered a general point if the expected return to self-employment is not simply identical to current income from employment, but rather there is some variation in the expected return. At the top tax threshold upside convexity for all tax payers goes from positive just below the threshold to zero just above it, because it is driven by what is happening to the marginal income tax rate in the base (i.e. in employment). To provide a simple illustration of this consider the

[16] Total convexity is constant within the income range £38,720 and £44,250, it takes value of 11, 14.5 and then 5.5 above and below these values, while upside convexity is constant within the income range £34,850 to the top tax threshold of £39,820, it takes value of 8, 18 and then 0 above and below these values.

example of two individuals, one with an income from employment that places them £1 below the point at which the marginal income tax rate changes and one with an income £1 above this threshold value. An individual with an income that places them just below the threshold point, will have positive 'upside' convexity, they will have a higher marginal income tax rate were they to become successful entrepreneurs (a positive value for convexity on the upside) and the same or lower tax rate if they are unsuccessful (whose value depends on the location of other thresholds and the variation in the expected returns from self-employment). In contrast, the individual with an income level that places them just above the top tax threshold already pays the highest marginal tax rate so their upside convexity must be zero, irrespective of how much they expect their incomes to rise if they were a successful entrepreneur.[17] Only those below the upper threshold are therefore subject to success taxes, while those tax payers above the top marginal tax threshold face no upside convexity.

Whilst the fall in upside convexity to zero is a general feature, the other values of convexity around the threshold will of course differ for every individual tax payer according to their expectations about the distribution of expected returns from self-employment. It might also be expected that values of tax convexity changes across time as their perceptions of the state of the world adjusts to new information and this affects their expectations about the risk and reward of starting their own enterprise. That however, highlights the importance of unobservable-time-varying changes in expectations in determining decisions about entrepreneurship and the construction of measures of tax convexity. These are of course factors that cannot be controlled for through the inclusion of individual specific time-invariant effects when using panels of micro data and therefore raises a question regarding the robustness of findings that rely on cross-section and time-variation in progressivity such as those by Gentry and Hubbard (2004, 2005). The RD approach, because it relies on cross-sectional differences, is arguably more robust to this point.

Finally in this section of the paper we describe what is happening to other relevant tax variables around the top tax threshold. In Figure 5 we show the average tax rate (left hand axis) paid by those in employment for an income window of between £35,000 to £45,000, alongside the differences between the average tax rate paid in employment and self-employment (right hand axis) around the top tax threshold for the same income window. We also show the probability weighted expected average tax in self-employment. To construct this last measure we use the same multiples of current income and probabilities used in the construction of the measure of tax convexity.

[INSERT FIGURE 5]

[17] This group has a positive value for downside convexity, because their tax rate will be lower if they are unsuccessful.

In all cases it is clear that these additional tax variables do not display the same discontinuity evident for upside and downside convexity in Figure 2. There is some variation in the tax variables within this income window however. The average tax rate is 27.3% for an employee with an income of £35,000 and 28.4% for an employee with an income of £45,000, while the expected average tax rate in self-employment is 25.3% and 27.4% at the same two points. The difference in the average tax rates paid in employment and self-employment across the same income window is 2.54 percentage points at £35,000 and 1.98 percentage points at £45,000. Those differences are smaller closer to the top tax threshold. At an income +/- £1,000 around the top tax threshold, which is the income window which is the basis of our preferred specifications, the values for the average tax rate are 26.9% and 27.1%, while the difference in the average tax rate paid in employment versus self-employment is 2.3 percentage points and 2.2 percentage points. The discontinuity is not therefore capturing differences in the average tax rate.

3. Data Description and Summary

[INSERT TABLE 2]

The data on self-employment used in this paper are drawn from the British Household Panel Survey (BHPS). This is an annual survey conducted in the United Kingdom with the central aim of collecting nationally representative longitudinal micro data on persons and households.[18] Households are sampled according to a stratified random cluster drawn from the population of British household postal addresses. Individuals who are aged over 16 years old within the household are administered the survey and assigned an individual identifier. The BHPS contains a fairly stable set of core questions covering the most essential areas such as health, education, labour market and socio-economic outcomes. Individuals are re-interviewed annually and at present annual waves of data are available to researchers from 1991. This database has been previously used to study the determinants of self-employment, including the effects of taxation by Parker (2003).

Our sample pools together information across the years 2000 to 2007. This contains the main variables of interest such as annual income and whether a person is employed or self-employed. Further control variables such as age, educational attainment, and a number of other variables are also taken from the BHPS data set. A summary of the main variables is provided in Table 2. Approximately 11% of individuals in the sample report being self-employed, while 1% of the sample enters self-employment during the period. There is relatively little cross time variation in this figure with a minimum entry rate of 0.87% (2005) and a maximum of 1.34% (2004) and no obvious trend in the data.

[18] It is similar to the SOEP for Germany and the PSID for the US.

4. Econometric Strategy

Our econometric strategy exploits exogenous variation in convexity created by sharp discontinuities in the MTR in the personal income threshold. For example, in 2007 workers whose income was below (the treatment group) the upper tax threshold of £39,825 paid a combined marginal tax rate of 23% marginal tax rate compared to the 41% rate paid by those in the higher tax bracket (the control group). Within a narrow neighbourhood of the threshold, whether a person's income lies just to the left or right of the cutoff is as good as randomly assigned (Lee, 2008; Lee and Lemieux, 2010). The evaluation problem consists of estimating the effect of the treatment (convexity) on the outcome variable (entry into self-employment or entrepreneurship). The key identification assumption that underlies the RD design is that the entry rate is a smooth (continuous) function of taxable income. Under this assumption the treatment effect can be retrieved by estimating the discontinuity in the local regression function at the point where the treatment indicator switches from 0 to 1.

Our estimation strategy uses both non-parametric and parametric methods. The non-parametric approach fits local polynomial regression functions either side of the higher rate threshold (HRT) and estimates the treatment effect as the jump in the rate of self-employment that occurs at the HRT. The choice of bandwidth in these applications is determined using the optimal bandwidth calculation described by Imbens and Kalyanaram (2009). Graphical plots of the relationship are based upon grouping the individual data into £500-wide bins and calculating the mean of the dependent variable within each bin. This bin width is selected based on optimal bin width tests reported in the Appendix. The parametric approach is conventional and involves estimating the equation

$$y_{it+1} = \alpha + \beta Treatment_{it} + \delta(\tilde{X}_{it}) + \gamma(\tilde{X} * Treatment_{it}) + Z_{it} + \varepsilon_{it} \quad (1)$$

where y_{it+1} is an outcome variable for individual i at time $t+1$ which takes a value of 1 if the individual enters self-employment/entrepreneurship in period $t+1$, and 0 otherwise. The variable of interest, $Treatment_{it}$, is a dummy equal to 0 if an individual's income (x_{it}) is greater than or equal to the HRT (x_0), 0 otherwise. That is

$$Treatment_{it} = \begin{cases} 1 \; if \; x_{it} < x_0, \\ 0 \; if \; x_{it} \geq x_{0.} \end{cases}$$

Our decision to model entry behavior in time period t+1 is motivated by the idea that the formation of business ideas takes time, but the results are unchanged when we use a contemporaneous dependent variable instead. The function $\delta(\tilde{X}_{it})$ includes polynomial expressions of the transformed assignment variable ($x_{it} - x_0$), where assignment is based on individual i's gross income in each year. Equation (1) also includes interactions between the transformed assignment variable and the treatment dummy to capture possible differences in the slope of the regression function either side of the threshold. In some specifications we also include a vector of individual characteristics (Z_{it}) that have been previously shown to affect the probability of entrepreneurship so as to reduce estimation bias. The equations are estimated using observations 'near' to the threshold. While we experiment with varying degrees of nearness, our preferred bandwidth (h) includes individuals with incomes £1,000 either side of the threshold to maximise precision.

[INSERT TABLE 3]

The validity of our identification strategy rests upon two important criteria. First, all other factors that may affect entrepreneurial outcomes must be continuous with respect to taxable income. If the treatment and control group systematically differ in terms of observable and unobservable characteristics, the control group will not represent the valid counterfactual and it will be impossible to disentangle whether the treatment effect is due to convexity or discontinuities in other covariates. Table 3 presents diagnostic tests that inspect how similar are the treatment and control groups. In Panel A we find no evidence of significant differences between the gender, age, marital status, incidence of parental self-employment, or educational attainment suggesting that the treatment and control groups around the HRT are indeed similar. We also test the comparability of the treatment and control groups around other tax thresholds. In Panel B of Table 3 we show that individuals around the allowance threshold are similar, but there is some indication that the treatment group around other marginal tax thresholds in the data. Those at 'threshold^1' are more likely to be male or to have had self-employed parents, while in Panel D none of the t-statistics are statistically significant at conventional levels suggesting that the groups either side of 'threshold^2' are broadly similar.

[INSERT FIGURE 6A, 6B] [INSERT TABLE 4]

A second important question is, are individuals able to influence the assignment variable, and if so, what is the nature of this control? Some individuals who are aware of the tax threshold may take steps to manipulate their income in order that they pay the lower MTR, leading to spurious results. Although we cannot test for strategic tax avoidance directly, because we only observe one observation of the assignment variable per individual per year, we follow convention in the literature and test whether the aggregate distribution of the assignment variable is continuous around the threshold (McCrary, 2008; Lee and Lemieux, 2010).[19] We begin by grouping the individual data into £500-wide bins and calculating the number of observations within each bin and plot a histogram and kernel density function of the

[19] Another possible source of discontinuity in the density of the assignment variable is selective attrition (Lee and Lemieux, 2010; DiNardo and Lee, 2004). In principle, if individuals within the high tax bracket are more likely to leave the sample (possibly to reduce monitoring by the tax authorities), there will be a discontinuity that could threaten the validity of the RD design. In that setting, testing for a discontinuity in the density is similar to testing for selective attrition (Lee and Lemieux, 2010). To investigate this issue we construct a binary variable equal to 1 if the individual exits the sample. A t-statistic (p-value) reveals no significant difference in the probability of exit between the treatment and control group: 0.0192 (0.44). We arrive at a similar conclusion when we use a probit model to regress the exit indicator on treatment status: coefficient [z-statistic] = -0.0840 [-0.77].

frequency distribution around the HRT in Figure 6A. The absence of a jump in the density of the assignment variable either just before the threshold, or at any other point of the distribution, indicates that sorting does not take place around the HRT. This is perhaps unsurprising as our sample comprises individuals in wage employment whose income is taxed at source.

In addition to the simple inspection of the data we also conduct more formal statistical tests for continuity in the assignment variable. Specifically, we use parametric and non-parametric RD methods to investigate whether a discontinuity in the frequency distribution coincides with the HRT. In the parametric case we estimate the equation

$$frequency_b = \alpha + \beta_1 Treatment_b + \beta_2 Assignment + \beta_3 (Assignment * Treatment_b) + \varepsilon_b \quad (2)$$

where $frequency_b$ is the number of observations in each £500-wide bin; $Treatment_b$ is a dummy variable equal to 1 if the bin is to the left of the HRT, 0 otherwise; $Assignment$ is the mid-point of each bin that assigns bins to treatment and control status; and ε_b is a stochastic error term. We also include only bins that lie within the range £10,000 above and below the HRT. The results of these tests are reported in Panel A in Table 4. Reassuringly, we do not find a statistically significant effect of treatment, indicating that sorting does not occur.[20] In Panel B of Table 4 we use fixed effects regression models to investigate whether individuals influence their treatment status by manipulating their income along other margins, for example by reducing the number of hours of overtime they work, or choosing part-time rather than full-time employment. None of the factors listed are found to be a significant determinant of treatment status which provides further support for the view that treatment status is as good as randomly assigned.[21]

5. Results

In this section we present estimates of the treatment effect obtained using non-parametric and parametric methods around the HRT, and study the robustness of our findings along several dimensions.

5.1 Non-Parametric Results

Figure 7 provides graphical evidence on the relationship between the rate of entry into self-employment and convexity (treatment). There is little evidence that points to an obvious discontinuity in the local polynomial regression functions plotted either side of the threshold, even if the local regression function is somewhat lower just to the right of the cut-off point within Figure 7. Rather, it would appear there is a high degree of variance in the probability of entry in each tax bracket.

[20] The results are unchanged when we use higher order polynomials in the parametric specification.
[21] Tests were also run to test for continuity at the other tax thresholds. Again, we did not find clustering to occur.

[INSERT FIGURE 7] [INSERT TABLE 5]

We investigate this more formally using non-parametric methods to estimate the size of the jump in the regression functions at the HRT. These estimates of the treatment effect of greater upside convexity of the personal income tax schedule are provided in Table 5. When we use only observations within a narrow neighbourhood of the HRT in column (1) of the table, we find that convexity has an unexpected positive influence on future decisions to become self-employed. While the magnitude of the effect varies between 0.014 and 0.018 percentage points according to which bandwidth is used, in all instances the coefficient is well below established levels of statistical significance. This provides initial evidence that there is no evidence for a discontinuity in the probability of entry due to treatment, a result we find is repeated throughout the analysis conducted in the paper. To put this differently, a conclusion that progressivity of the UK income tax schedule acts as 'success tax' on entrepreneurs is not supported by the data.

Given that the regression discontinuity design tends to rely on a small sample size, there is an inherent trade-off between the precision and efficiency of its estimates. A concern could be that we find upside tax convexity to be an insignificant determinant of entry because of the limited number of data points we use in the estimation of the regression which inflates the standard errors. To investigate whether this is the case we expand the sample window so that it encompasses observations further from the threshold. In regressions 3, 5 and 7 we report the treatment effect based on a sample that includes observations within a range of £2,000, £5,000 and £10,000 either side of the threshold. As the size of the income band around the threshold, and the number of observations increases across the table we find if anything that the coefficient estimate on the convexity variable (*treatment*) becomes closer to zero and is at no point statistically significant.

An alternative explanation for the general insignificance observed in Table 5 could be that the coefficient on the treatment dummy is biased downwards because of the omission of other relevant factors, and that these influences are negatively correlated with treatment status. Given the findings presented previously in Table 3 this seems to be an unlikely concern. Further support for this view can be found in columns (2), (4), (6), and (8) of Table 5, where we find that the treatment effect remains insignificant when additional characteristics of the individual, including their age and marital status, are included in the estimated regression.[22]

[INSERT TABLE 6]

5.2 Parametric Results

[22] We also sought to exploit information within the BHPS on the industry in which an individual is currently employed so as to restrict the sample to include only individuals currently employed in occupations where self-employment is common (e.g. dentists, builders, taxi drivers etc.), and convexity would be expected to have a correspondingly greater influence on transitions into self-employment. However, there were not sufficient data points to run the regressions.

In order to ensure our findings are not simply due to the type of estimation strategy used, we also parametrically estimate equation (1). Again we experiment with different bandwidths. Given the narrow data range used for estimation we cannot afford to be as agnostic about the functional form of the regression line when using the RD design compared with other estimators. Correct specification of the functional form of the regression function is paramount because this may affect the size of the discontinuity in the local polynomial regressions we estimate either side of the threshold (Lee and Lemieux, 2010; Angrist and Pischke, 2009). Moreover, based on the evidence in Figure 7, we therefore include interactions between the transformed assignment variable and the treatment indicator to obviate the bias that would arise by imposing the same functional form either side of the threshold. The results of these tests are provided in Table 6. From this exercise we conclude that our results are robust – the parametric results also indicate that success taxes are an insignificant determinant of entry into self-employment.[23]

[INSERT TABLE 7]

5.3 Omitted Tax Variables

Provided that strategic behaviour is absent from the data, the regression discontinuity design provides causal inference as treatment status is randomly assigned and therefore uncorrelated with potential confounding omitted variables contained within the error term. In practice researchers using RD designs include other covariates because doing so can reduce the sampling variability of the estimator. Arguably the greatest potential for this occurs when one of the baseline covariates is a pre-random-assignment observation on the dependent variable, which may likely be highly correlated with the post-assignment outcome variable of interest (Lee and Lemieux, 2010). In this section we consider the omitted control variables to include other tax variables (Table 7) or alternative estimates of tax convexity.

The list of tax variables that the literature on taxation and entrepreneurship has considered to be potentially important is somewhat long and so in this section we focus on a variants of the probability weighted expected average tax rates that Bacher and Brülhart (2013) show affect the results for tax convexity in corporate tax rates and new firm entry. Following evidence from Robson and Wren (1999), Schuetze (2000), Bruce (2000) and Cullen and Gordon (2002), and Fairlie and Mayer (1999) we include a measure of the average tax rate paid in employment as well as the difference in the average tax rate in employment and self-employment at the current wage income of the employee. We also include a measure of the probability weighted expected average tax rate in self-employment, where we use the same probabilities and income multiples that we have used in the construction of our measure of tax convexity and were used by Gentry

[23] Additional tests were run using higher order polynomials and asymmetric polynomials. The results were unchanged.

and Hubbard (2000, 2004, 2005).[24] The final measure of average tax rates we include uses only the expected successful outcomes of entrepreneurship in constructing the expected average tax rate in self-employment.

In column (1) in Table 7 we report the effect of including the average tax rate paid in employment; in column (2) the we include a measure of the difference in the average tax rate paid in employment and self-employment; column (3) the probability weighted expected average tax rate in self-employment; and in column (4) the probability weighted average tax rates in successful states of self-employment. Throughout the results reported in Table 7 we continue to find that progressivity of the income tax schedule has no effect on self-employment. However there does appear to be an effect from the tax variables measuring the difference in the tax rates in employment and self-employment at the current income of the employee. The results for this variable in column 2 suggest that the greater is this difference, the more likely it is a person will become self-employed in the next time period.

The significance of this variable combined with the insignificance of the success tax variable lends support to the Bacher and Brülhart (2013) argument that omitted tax variables might explain the results in Gentry and Hubbard (2000, 2004, 2005). Given the evidence from Figure 5 of limited variation in this measure in the narrow income window that we study, why this variable has a significant effect on the decision to become self-employed is somewhat of a puzzle. It seems to indicate a very-strong non-linearity in the effects of this tax variable. A theoretical explanation for why that occurs is outside of the scope of the paper. We include the difference in the average tax rate paid in employment and self-employment as a control variable throughout the rest of the analysis, although we note that we do not always find it to be a robust determinant of entry into self-employment.

In Table 7 column 5 we also test the robustness of our findings to including the convexity measure constructed by Wen and Gordon (2013), following as closely as possible the methodology outlined in that paper.[25] To construct this convexity variable we first restrict the sample to observations of individuals in self-employment and regress income on the average tax rate difference, marriage indicator, gender, age, parental self-employment, the university degree dummy, and year dummies and save the estimated coefficients. Next, for each individual in *employment* we construct expected earnings in self-employment by taking 1,000 draws from a

---

[24] Bacher and Brülhart (2013) criticise the work of Gentry and Hubbard (2000, 2005) for not including this variable (see footnote 10, page 133). We also used a measure of the probability weighted difference in the expected average tax rate in self-employment and the employment. This variable was found to be statistically insignificant.

[25] Cullen and Gordon (2007) extend Gentry and Hubbard (2000) to include the effect of the difference between personal and corporate tax rates on risk taking. They argue that successful entrepreneurs can avoid high personal income tax rates by incorporating. As both the marginal and average corporate tax rates in the UK are considerably lower than the personal income tax rate at the top personal income tax threshold, successful enterprises might choose to incorporate in order to minimise their overall tax bill Crawford and Freedman (2010). However (2004) show that changes to the corporate tax rate in the UK over the sample period have affected the decision of the self-employed to incorporate, the exit margin, and there is little suggestion it has affected the decision to leave employment for self-employment, the entry margin that we study here.

normal distribution with a conditional mean and variance equal to the respective values from the earnings of the self-employed in our data set. The expected earnings of entrepreneurship for each individual in employment is then calculated as the sum of their expected earnings and observed characteristics.[26] Convexity^{WG} is then calculated based on the difference between the individual's observed income from employment and their expected earnings in entrepreneurship. Unlike the success tax construction of convexity we note there are no sharp discontinuities in the values around the upper tax threshold.

As reported in Table 7, we do not find the Wen and Gordon convexity measure to be a statistically significant determinant of entry within the narrow sample window that we exploit.[27] Its inclusion also leaves the previous findings for the treatment indicator unaffected.

[INSERT TABLE 8]

5.4 Long Run Outcomes

We might conclude from the above analysis that the convexity of the tax schedule neither encourages nor discourages employees paying the higher rate tax to switch into self-employment in the next year. But what if those decisions take longer than a single year to make; i.e. the period between treatment and the gestation of an idea into a business takes several years? In column (1) in Table 8 we explore whether those employees paying the higher rate marginal tax rate in 2000 were more likely to become self-employed in *any* year up to 2007. As is evident from this table we find no statistically significant effect of facing higher upside convexity and becoming self-employed in any of the next seven years. Success taxes have no long-run effects. This result is robust to the addition of those that are subject to success taxes in either 2000 or 2001 [column (2)] or any year up to and including 2000 to 2002 [column (3)]. The results are unchanged when we include higher order polynomials in the estimating equation.

[INSERT TABLE 9]

5.5 Other Thresholds

As a final exercise in the robustness of our main findings we use an RD approach around the other thresholds in the combined marginal income tax rate that occur in the UK. Column (1) in Table 9 explores the point at which the marginal income tax rate jumps from 0% to 11% (when income reached £5,200 in 2007), column (2) from 21% to 33% (an income of £7,455 in 2007), and column (3) from 33% down to 23% (an income of £34,840 in 2007). The values of upside (downside) convexity £1 above and below the threshold for the year 2007 are 25.2 and 4.2 (0

[26] That is $E(x)_{it} = \frac{1}{N}\sum_{j=1}^{1000}(\mu)_i + \hat{\alpha}_1(ATAX)_{it} + \hat{\alpha}_2(Married)_{it} + \hat{\alpha}_3(Male)_{it} + \hat{\alpha}_4(Parent\ self-employed)_{it} + \hat{\alpha}_5(Degree)_{it}$. Where the parameters $j \in (\hat{\alpha}_1,\ \hat{\alpha}_2, \hat{\alpha}_3, \hat{\alpha}_4, \hat{\alpha}_5)$ are the coefficient estimates from the self-employment regression.

[27] When we include this measure of convexity within a larger sample i.e. one that does not restrict observations to be in the location of the upper tax threshold, we find the expected significant effect for this variable.

and 11) in column (1); 9 and -3 (6.6 and 18.6) in column (2); and 8 and 18 (3 and -7) in column (3). As is evident from Figure 2 there are two points at which the combined marginal income tax changes for very low incomes. In column (1) we examine the threshold at which the marginal income tax rate jumps from 0% to 11%, which was £5,200 in 2007, although the results are likely to be affected by the next threshold which lies very close by, £5,225 in 2007. The expected signs for treatment are positive in regressions 1 and 2, as upside convexity falls to the right of the threshold, and negative in regression 3, as upside convexity is higher to the right of the threshold point. However, in all three instances the sign on the estimated coefficient is the opposite of the predicted sign. The evidence from Table 10 indicates that there are no significant differences in the probability of becoming self-employed for employees that face different marginal tax rates and therefore different degrees of progressivity of the tax schedule. Success taxes seem to be unimportant throughout the income tax schedule in the UK.

6. Conclusions

Entrepreneurship is often associated with having many beneficial effects on the economy, including on employment and innovation. For this reason there has been a long-established interest from policy makers and academics regarding how policy might be designed to encourage these positive effects. Within this debate, the effects of progressivity of the income tax policy schedule has received less attention compared to those of the level and changes to average and marginal tax rates. One exception to this has been the work of Gentry and Hubbard (2000, 2004, 2005) who provide strong empirical evidence that greater convexity is negatively correlated with entry into self-employment using US data. Progressivity taxes success, discouraging entrepreneurship.

In this paper, motivated by the evidence of these success taxes for the US, we exploit the sharp discontinuity that occurs in their value at the top tax threshold to explore the effects of progressivity on self-employment in UK data. The success tax argument would predict that individuals with an income that is just below the top tax threshold face a greater disincentive to move into self-employment compared to those with an income just above the top tax threshold. We can find no such evidence of an effect within the UK data, a result that is robust to a number of checks including examining this question over longer time horizons, addition of other controls, or using other threshold points.

The absence of evidence might be viewed either in a positive or negative manner by policy makers. In the absence of evidence for an effect from this aspect of progressivity they are freed from a need to consider this when trying to design tax policy. Alternatively, they are denied a policy tool that might otherwise have been used to encourage entrepreneurs and entrepreneurship.

Why might our results differ compared to those of Gentry and Hubbard (2004, 2005) and Wen and Gordon (2013)? With regard to the latter there is an obvious difference with the measure of tax convexity that is used. With regard to the former, with which our paper is more closely aligned, the difference in context, that we use UK data and they the US, might provide one possibility, while the difference in methodology offers an alternative explanation. There may be a number of aspects to the effect of the difference in methodology. Firstly, Bacher and Brülhart (2013) criticise Gentry and Hubbard (2000, 2004, 2005) for their failure to control for other tax variables in their regressions and it may be possible that they capture the effects of these taxes rather than progressivity. We find some support for this. We find that the difference in the average tax rate paid in employment and self-employment does help to predict future transitions into self-employment and that this occurs at points close to the top tax threshold, where tax convexity also differs. It might alternatively be that by relying on cross-time changes in tax rates Gentry and Hubbard (2000, 2004, 2005) capture unobservable time-varying individual-specific factors. An individual's response to *changes* in tax rates may occur because of changes in convexity, or because of a change in their expectations about the returns to entrepreneurship, which are likely to be dependent upon their expectations of the state of the world now, and in the future. These are factors that may be correlated with variables that determine the change in the income tax rate. The reliance of the regression discontinuity approach on differences between individuals within a year arguably does a better job of holding constant unobservable factors such as expectations about the current and future state of the economy. Either way the differences between the UK and US results suggest that further investigation of this is issue is warranted.

References

Angrist, J.D. and Pischke, J.S. (2009). 'Mostly Harmless Econometrics', 1st ed., Oxford: Princeton University Press.

Bacher, H.U. and Brülhart, M. (2013). 'Progressive Taxes and Firm Births', *International Tax and Public Finance*, Vol. 20, pp. 129-168.

Bartelsman, E, Scarpetta, S and Schivardi, F. (2005). 'Comparative Analysis of Firm Demographics and Survival: Evidence from Micro-Level Sources in OECD countries', *Industrial and Corporate Change*, Vol. 14, pp. 365-391.

Baumol, W.J. (2002). 'The Free-Market Innovation Machine: Analyzing the Growth Miracle of Capitalism', Princeton University Press.

Blau. D. M., (1987). 'A Time-Series Analysis of Self-Employment in the United States', *The Journal of Political Economy*, Vol. 95, pp. 445-467

Bruce, D. (2000) 'Effects of the U.S. Tax System on Transitions into Self-employment', *Labour Economics*, Vol. 7, pp. 547-575.

Bruce, D. and Deskins. J. (2012). 'Can State Tax Policies be Used to Promote Entrepreneurial Activity?', *Small Business Economics*, Vol. 38, pp. 375-397.

Bruce, D. and Mohsin. M. (2006). 'Tax Policy and Entrepreneurship: New Time Series Evidence', *Small Business Economics*, Vol. 26, pp. 409-425.

Carroll, R., Holtz-Eakin, D., Rider, M. and Rosen, H.S. (1998). 'Entrepreneurs, Income Taxes, and Investment', NBER Working Paper, No. 6374.

Crawford, C. and Freedman, J. (2010), "Small Business Taxation", in *Dimensions of Tax Design: the Mirrlees Review*, J. Mirrlees, S. Adam, T. Besley, R. Blundell, S. Bond, R. Chote, M. Gammie, P. Johnson, G. Myles and J. Poterba (eds.) Oxford University Press, Book Chapters

Cullen, J., & Gordon, R. (2002). 'Taxes and Entrepreneurial Activity: Theory and Evidence for the US', NBER Working Paper No 9015.

Cullen, J., and Gordon, R.H. (2007). 'Taxes and Entrepreneurial Risk-Taking: Theory and Evidence for the U.S.', *Journal of Public Economics*, Vol. 91, pp. 1479-1505.

Da Rin, M., Di Giacomo, M., and Sembenelli, A. (2011). 'Entrepreneurship, Firm Entry, and the Taxation of Corporate Income: Evidence from Europe', *Journal of Public Economics*, Vol. 95, pp. 1048-1066.

DiNardo, J., and Lee, D. S. (2004). 'Economic Impacts of New Unionization on Private Sector Employers: 1984–2001', *The Quarterly Journal of Economics*, Vol. 119(4), pp. 1383-1441.

Disney, R., Haskel, J. and Heden, Y. (2003). 'Restructuring and Productivity Growth in UK Manufacturing', *Economic Journal*, Vol. 113, pp. 666-694.

Djankov, S., Ganser, T., McLiesh, C., Ramalho, R. and Schleifer, A. (2010). 'The Effect of Corporate Taxes on Investment and Entrepreneurship', *American Economic Journal: Macroeconomics*, Vol. 2, pp. 31-64.

Fairlie, R. W., & Meyer, B. D. (2000). 'Trends in Self-Employment among White and Black Men During the Twentieth Century', *Journal of Human Research*, Vol. 35, pp. 643–669.

Ferede, E. (2013). 'Tax Progressivity and Self-Employment: Evidence from Canadian Provinces', *Small Business Economics,* Vol.40 ,pp. 141-153.

Fölster, S. (2002). 'Do Lower Taxes Stimulate Self-Employment?', *Small Business Economics*, Vol 19, pp. 135–145.

Gentry, W., and Hubbard, R.G. (2000). 'Tax Policy and Entrepreneurial Entry', *American Economic Review*, Vol. 90, pp. 283-287.

Gentry, W., and Hubbard, R.G. (2004). 'Tax Policy and Entry into Entrepreneurship,' mimeo. Colombia University.

Gentry, W., and Hubbard, R.G. (2005). 'Success Taxes, Entrepreneurial Entry, and Innovation', with R. Glenn Hubbard, in Adam B. Jaffe, Joshua Lerner and Scott Stern, eds., Innovation Policy and the Economy, volume 5, MIT Press, pp. 87-108.

Graham. J.R. and Rogers, D.A. (2002). 'Do Firms Hedge in Response to Tax Incentives', *The Journal of Finance*, Vol. LVII, pp. 815-839.

Graham. J.R. and Smith, C.W. (1999). 'Tax Incentives to Hedge'. *Journal of Finance*, Vol. LIV, pp. 2241-2262.

HMRC (2004), *Regulatory Impact Assessment (RIA)—Corporation Tax: The Non-Corporate Distribution Rate*, London: HMRC.

Imbens, G. and Kalyanaraman, K. (2009). 'Optimal Bandwidth Choice for the Regression Discontinuity Estimator', NBER Working Papers No. 14726.

Lee, D. (2008) 'Randomized Experiments from Non-Random Selection in U.S. House Elections', *Journal of Econometrics*, Vol. 142, pp. 675-697.

Lee, D. and Lemieux, T. (2010). 'Regression Discontinuity Designs in Economics', *Journal of Economic Literature*, Vol. 48, pp. 281-355.

Long, J. E. (1982). 'The Taxation of Self-Employment', *National Tax Journal*, Vol. 35, pp. 31-42.

McCrary, J. (2008) 'Manipulation of the Running Variable in the Regression Discontinuity Design: A Density Test', *Journal of Econometrics*, Vol. 142, pp. 698-714.

Musgrave, R. A., and Musgrave, P. B. (1989). 'Public Finance in Theory and Practice', London: McGraw-Hill.

Musgrave, R.A. and Tun Thin (1948). 'Income Tax Progression 1929-48', *The Journal of Political Economy*, Vol. 56, pp. 498-514.

OECD (2001). 'Drivers of Growth: Information Technology, Innovation and Entrepreneurship', *Science, technology and industry outlook*, Paris, OECD.

Parker, S. C. (2003). 'Does Tax Evasion affect Occupational Choice?', *Oxford Bulletin of Economics and Statistics*, Vol. 65(3), pp. 379–394.

Robson, M., Wren, C. (1999). 'Marginal and Average Tax Rates and the Incentive for Self-Employment', *Southern Economic Journal*, Vol. 65, pp. 757– 773.

Saez, E. (2010). 'Do Taxpayers Bunch at Kink Points?', *American Economic Journal: Economic Policy*, Vol. 2, pp. 180-212.

Sarkar, S. (2008). 'Can Tax Convexity be Ignored in Corporate Finance Decisions?' *Journal of Banking and Finance*, Vol. 32, pp. 1310-1321.

Sarkar, S. and Goukasian, L. (2008). 'The Effects of Tax Convexity on Corporate Investment and Tax Burdens', *Journal of Public Economic Theory*, Vol. 8, pp. 293-320.

Schuetze, H. (2000). 'Taxes, Economic Conditions and Recent Trends in Male Self-Employment: A Canada-US Comparison', *Labour Economics*, Vol. 7,pp. 507–544.

Wen, J-F. and Gordon, D. V. (2013). 'An Empirical Model of Tax Convexity and Self-Employment', *Review of Economics and Statistics*, (forthcoming).

Tables

Table 1
UK Combined Personal Income & NIC Tax Structure

The table presents information on the marginal tax rate levied in each tax bracket of the UK personal income tax system for individuals in employment (Panel A) and self-employment (Panel B). The allowance threshold is the level of income below which the marginal income tax rate is 0%. Individuals who earn an income greater than or equal to the allowance but less than the next income threshold, Threshold^1 in the case of employed individuals, pay a marginal tax rate of τ^1. Where an individual's income is greater than or equal to Threshold^1, but less than Threshold^2, the marginal tax rate is set at τ^2. Individuals who earn an income of at least Threshold^2 pay a marginal tax rate of τ^3. Threshold^4 corresponds to the higher rate threshold (HRT).

Year	Allowance	τ^1	Threshold^1	τ^2	Threshold^2	τ^3	Threshold^3	τ^4	Threshold^4	τ^5
					A: Employees					
2000	3,952	10	4,385	20	5,905	32	27,820	22	32,785	40
2001	4,524	10	4,535	20	6,415	32	29,900	22	33,935	40
2002	4,615	10	4,628	20	6,535	32	30,420	22	34,515	40
2003	4,615	10	4,628	21	6,575	33	30,940	23	35,115	41
2004	4,745	21			6,765	33	31,720	23	36,145	41
2005	4,888	11	4,895	21	6,985	33	32,760	23	37,295	41
2006	5,035	10	5,044	21	7,185	33	33,540	23	38,335	41
2007	5,200	11	5,225	21	7,455	33	34,840	23	39,825	41
					B: Self-Employed					
2000	4,385	17			5,905	29	27,820	22	32,785	40
2001	4,385	17			6,415	29	29,900	22	33,935	40
2002	4,615	17			6,535	29	30,420	22	34,515	40
2003	4,615	18			6,575	30	30,940	23	35,115	41
2004	4,745	18			6,765	30	31,720	23	36,145	41
2005	4,895	18			6,985	30	32,760	23	37,295	41
2006	5,035	18			7,185	30	33,540	23	38,335	41
2007	5,225	18			7,455	30	34,840	23	39,825	41

Table 2
Summary Statistics

The table presents summary statistics for the sample window £1,000 either side of the higher rate threshold. The sample includes all economically active individuals aged between 16 and 65.

Variable	*N*	Mean	Std. Dev.	Min	Max
Entry	813	0.0172	0.1302	0	1
Total convexity	813	14.0824	1.1352	11	14.5
Upside convexity	813	9.9188	8.9585	0	18
Downside convexity	813	4.1636	9.3969	-7	14.5
Average tax rate in employment	813	26.53	0.4596	25.79	27.18
Difference in average tax rate employment and self-employment	813	2.24	0.0752	1.76	2.35
Expected average tax rate in self-employment	813	24.29	0.4858	23.51	24.99
Expected average tax rate in successful self-employment	813	29.45	0.5279	28.46	30.14
Income	813	35,822	2,226	31,800	40,603
Married	813	0.7724	0.4195	0	1
Male	813	0.7257	0.4464	0	1
Age	813	42.42	8.9929	16	65
Parent self-employed	813	0.0086	0.0924	0	1
University degree	813	0.3112	0.4633	0	1
Hours worked per week	812	38.63	7.86	0	98
Overtime hours	803	6.57	7.12	0	40
Paid overtime hours worked	813	1.31	4.18	0	35
Second job	811	0.05	0.22	0	1
Bonuses	813	1,040.92	3871	0	53,000
Tenure	811	5.41	6.64	0	34
Full-time job (30 hours +)	811	0.98	0.13	0	1

Table 3
Testing for Discontinuities in Other Covariates & Manipulation of Treatment Status

The table provides diagnostic tests that investigate whether there are jumps in other covariates coinciding with the discontinuity in convexity at various tax thresholds. We use *t*-tests to investigate if there are statistically significant differences in the covariates among the treatment and control groups each side of the threshold. In all cases the control group comprise individuals with an income up to £1,000 more than the threshold, whereas the treatment group constitute all individuals with an income between at least the threshold and £1,000 below it. 'Treatment' ('Control') is the mean value of each covariate among the treatment (control) group. Panel A presents the tests for the sample around the HRT. Panel B presents the tests for the sample around the allowance threshold (where τ^1 takes effect). Panel C presents the tests for the sample around Threshold[1] (where τ^2 takes effect). Panel D presents the tests for the sample around Threshold[2] (where τ^3 takes effect). The *t*-statistic reports the difference between the mean values of covariate *j* within the treatment and control group.

	Panel A: Higher rate threshold			Panel B: Allowance threshold			Panel C: Threshold[1]			Panel D: Threshold[2]		
Variable	Control	Treatment	*t*-statistic	Control	Treatment	*t*-statistic	Control	Treatment	*t*-statistic	Control	Treatment	*t*-statistic
Married	0.7616	0.7813	-0.66	0.6085	0.6270	-1.07	0.6275	0.6355	-0.48	0.7150	0.7480	-1.29
Male	0.7041	0.7433	-1.25	0.1414	0.1549	-1.08	0.1765	0.1416	2.75**	0.6938	0.6803	0.50
Age	42.32	42.50	-0.28	39.39	40.05	-1.31	39.13	39.76	-1.34	41.45	41.76	-0.57
Parent self-employed	0.0027	0.0134	-1.63	0.0104	0.0117	-0.34	0.0196	0.0102	2.21*	0.0035	0.0126	-1.72+
University degree	0.3068	0.3147	-0.24	0.0547	0.0449	1.26	0.0521	0.0596	-0.94	0.2973	0.2646	1.26

Table 4
Test for Continuity in the Assignment Variable at the Threshold

Panel A of the table presents parametric and non-parametric tests of the hypothesis that there exists a discontinuity in the frequency of the assignment variable at the cutoff point. Frequency is counted by summing the number of individuals in each £500-wide bin. h = £10,000 in both regressions. The results in column (1) are estimated using the optimal bandwidth calculation by Imbens and Kalyanaram (2009) to minimize the mean square error. z-statistics reported in parentheses in column (1), heteroskedasticity robust t-statistics reported in parentheses in column (2). ** indicates significance at the 1% level. Panel B reports the results from regressions that test for the exogeneity of treatment by estimating the equation $Treatment_{i,t} = \alpha + \beta X_{i,t} + \gamma_i + \gamma_t + \varepsilon_{it}$, where X_{it} is an explanatory variable, γ_i denote individual fixed effects, γ_t are year fixed effects and ε_{it} is a stochastic error term. **, * and + indicate statistical significance at the 1%, 5% and 10% level.

Panel A: Continuity tests	(1) Non-Parametric	(2) Parametric
Treatment (success taxes)	18.7894	-20.4737
	(1.20)	(-0.71)
Assignment		-38.4466**
		(-6.33)
Assignment * Treatment		-25.4316**
		(-4.10)
N	40	40

Panel B: Manipulation of treatment status

$X_{i,t}$	β	t-stat	N	R^2
Hours worked per week	0.0064	1.00	812	0.68
Overtime hours	0.0074	1.06	803	0.68
Paid overtime hours worked	0.0300	1.27	813	0.71
Second job	-0.1351	-0.74	811	0.68
Bonuses	0.0001	1.49	813	0.68
Tenure	0.0080	-0.87	811	0.68
Full-time job (30 hours +)	0.0251	0.08	811	0.68

Table 5
Non-Parametric Estimates

The table presents non-parametric regression discontinuity estimates of the treatment effect. Each cell represents a single regression. The coefficients reported in the first row of the table are estimated using the optimal bandwidth, calculated using the methodology outlined by Imbens and Kalyanaram (2009) to minimize the mean square error. $\text{Bandwidth}_{0.5}$ and Bandwidth_{2} estimate the treatment effect using half and twice the optimal bandwidth respectively. *h* denotes the width of the sample window around the higher rate threshold. For example, *h* = £1,000 indicates that the sample comprises observations within a range £1,000 above and below the threshold. The dependent variable in all regressions is $enter_{i,t+1}$. Covariates are: married, age, marital status, parental self-employment, and the university degree indicator. *z*-statistics reported in parentheses. **, * and + indicate significance at the 1%, 5% and 10% level.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	h = £1,000		*h* = £2,000		*h* = £5,000		*h* = £10,000	
Bandwidth_{1}	0.0181	0.0140	0.0134	0.0118	0.0138	0.0120	0.0141	0.0122
	(1.12)	(0.97)	(1.03)	(0.95)	(1.04)	(0.94)	(1.05)	(0.94)
$\text{Bandwidth}_{0.5}$	0.0143	0.0096	0.0188	0.0150	0.0185	0.0146	0.0178	0.0139
	(0.96)	(0.68)	(1.32)	(1.12)	(1.28)	(1.09)	(1.22)	(1.03)
Bandwidth_{2}	0.0180	0.0152	0.0162	0.0151	0.0062	0.0055	0.0068	0.0061
	(1.23)	(1.07)	(1.32)	(1.26)	(0.59)	(0.54)	(0.65)	(0.59)
Covariates		√		√		√		√
N	813	813	1,617	1,617	4,063	4,063	9,319	9,319

Table 6
Parametric Estimates

The table presents parametric regression discontinuity estimates of the treatment effect around the higher tax rate threshold. Each cell represents a single regression. We choose not to report the coefficient estimates for the polynomial functions, the interactions between the polynomial functions and treatment dummy, and the covariates for the sake of parsimony (the regression output is available on request). *h* denotes the width of the sample window around the higher rate threshold. For example, *h* = £1,000 indicates that the sample comprises observations within a range £1,000 above and below the threshold. The dependent variable in all regressions is $enter_{i,t+1}$. Covariates are: married, age, marital status, parental self-employment, and the university degree indicator. Heteroskedasticity robust *t*-statistics reported in parentheses. *p*-values from the goodness-of-fit test are reported in square brackets. The goodness-of-fit test is computed by jointly testing the significance of a set of bin dummies included as additional regressors in the model with the associated *p*-values reported. The bin width used to construct the bin dummies is £500. **, * and + indicate significance at the 1%, 5% and 10% level.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Bandwidth	*h* = £1,000		*h* = £2,000		*h* = £5,000		*h* = £10,000	
Polynomial of order:								
One	-0.0188	-0.0204	0.0168	0.0185	-0.0077	-0.0087	-0.0015	-0.0014
	(-0.43)	(-0.66)	(0.91)	(1.27)	(-0.95)	(-1.07)	(-0.31)	(-0.26)
	[0.50]	[0.34]	[0.81]	[0.85]	[0.12]	[0.29]	[0.12]	[0.32]
Two	-0.0143	-0.0183	-0.0191	-0.0178	0.0086	0.0107	-0.0048	-0.0054
	(-0.29)	(-0.47)	(-0.52)	(-0.60)	(0.61)	(0.81)	(-0.59)	(-0.64)
	[0.51]	[0.34]	[0.44]	[0.47]	[0.62]	[0.73]	[0.14]	[0.33]
Three	-0.0212	-0.0210	-0.0184	-0.0204	-0.0019	-0.0005	-0.0062	-0.0069
	(-0.37)	(-0.53)	(-0.41)	(-0.51)	(-0.08)	(-0.03)	(-0.54)	(-0.58)
	[0.68]	[0.30]	[0.54]	[0.55]	[0.53]	[0.63]	[0.24]	[0.44]
Four	-0.0135	-0.0142	-0.0015	-0.0023	0.0090	0.0102	0.0080	0.0092
	(-0.20)	(-0.40)	(-0.03)	(-0.06)	(0.27)	(0.32)	(0.49)	(0.58)
	[0.62]	[0.23]	[0.37]	[0.26]	[0.73]	[0.81]	[0.58]	[0.75]
Covariates		√		√		√		√
N	813	813	1,617	1,617	4,063	4,063	9,319	9,319

Table 7
Other Tax Variables

The table presents parametric regression discontinuity estimates of the treatment effect around the higher tax rate threshold but incorporates other relevant tax variables in the specification. The dependent variable in all regressions is $enter_{i,t+1}$. In all columns the sample is restricted to a range £1,000 either side of the cutoff, i.e. h = £1,000. Heteroskedasticity robust t-statistics reported in parentheses. The goodness-of-fit test is computed by jointly testing the significance of a set of bin dummies included as additional regressors in the model with the associated p-values reported. The bin width used to construct the bin dummies is £500. **, * and + indicate significance at the 1%, 5% and 10% level.

	(1)	(2)	(3)	(4)	(5)
Bandwidth: h = £1,000					
Treatment (success taxes)	-0.0183	-0.0193	-0.0183	-0.0183	-0.1595
	(-0.58)	(-0.61)	(-0.58)	(-0.42)	(-0.90)
Assignment	-0.0409	-0.0379	-0.0401	-0.0371	-0.0396
	(-1.01)	(-0.92)	(-0.99)	(-0.75)	(-0.98)
Assignment * Treatment	-0.0875+	-0.0905+	-0.0870+	-0.0825	-0.0821
	(-1.70)	(-1.73)	(-1.69)	(-1.14)	(-1.62)
Average tax rate in employment	-0.0070				
	(-0.68)				
Difference in average tax rate employment and self-employment		0.0676*			
		(2.00)			
Expected average tax rate in self-employment			-0.0077		
			(-0.78)		
Expected average tax rate in successful self-employment				-0.0079	
				(-0.79)	
Convexity^{WG}					-0.0078
					(-0.83)
Goodness-of-fit	0.38	0.36	0.39	0.55	0.39
N	813	813	813	813	813

Table 8
Long-Run Effects

The table presents parametric regression discontinuity estimates of the long-run treatment effect. In all columns the sample is restricted to a range £1,000 either side of the cutoff, i.e. h = £1,000. In column (1) the dependent variable is equal to 1 if an individual observed in 2000 enters self-employment in any year between 2001 and 2007, 0 otherwise. In column (2) the dependent variable is equal to 1 if an individual observed in 2000 and 2001 enters self-employment in any year between 2002 and 2007, 0 otherwise. In column (3) the dependent variable is equal to 1 if an individual observed in 2000, 2001 and 2002 enters self-employment in any year between 2003 and 2007, 0 otherwise. Column (1) uses observations from 2000 only. Column (2) uses observations from 2000 and 2001 only. Column (3) uses observations from 2000, 2001 and 2002 only. Heteroskedasticity robust t-statistics reported in parentheses. The goodness-of-fit test is computed by jointly testing the significance of a set of bin dummies included as additional regressors in the model with the associated p-values reported. The bin width used to construct the bin dummies is £500. **, * and + indicate significance at the 1%, 5% and 10% level.

	(1)	(2)	(3)
Tax Treatment in year(s)	2000	2000-01	2000-02
Bandwidth: h = £1,000			
Treatment (success taxes)	-0.0920	-0.1010	-0.0359
	(-1.12)	(-1.37)	(-0.34)
Assignment	-7.5980	-0.1960	-0.1478
	(-1.41)	(-1.42)	(-1.02)
Assignment * Treatment	-0.4600	-0.3519*	-0.1519
	(-1.45)	(-2.01)	(-1.13)
Difference in average tax rate	-104.3427	-0.4642	-0.4487
employment and self-employment	(-1.41)	(-0.46)	(-0.58)
Goodness-of-fit	0.33	0.23	0.93
N	97	198	312

Table 9
Alternative Tax Thresholds

The table presents parametric regression discontinuity estimates of the treatment effect around the other tax thresholds listed in Table 1. In all columns the sample is restricted to a range £1,000 either side of the cutoff, i.e. h = £1,000. The dependent variable in all regressions is $enter_{i,t+1}$. In all regressions the treatment (control) group are observations to the left (right) of the threshold. Heteroskedasticity robust t-statistics reported in parentheses. The goodness-of-fit test is computed by jointly testing the significance of a set of bin dummies included as additional regressors in the model with the associated p-values reported. The bin width used to construct the bin dummies is £500. **, * and + indicate significance at the 1%, 5% and 10% level.

Tax Threshold	(1) Allowance	(2) Threshold[1]	(3) Threshold[2]
Treatment	0.0064	0.0036	-0.0099
	(0.83)	(0.47)	(-0.70)
Assignment	-0.0080	-0.0252	-0.0059
	(-0.43)	(-1.16)	(-0.25)
Assignment * Treatment	-0.0262	-0.0202	0.0448
	(-1.01)	(-0.78)	(1.27)
Difference in average tax rate employment and self-employment	0.0060	-0.0053	0.1416
	(0.71)	(-0.63)	(1.32)
Goodness-of-fit	0.92	0.23	0.29
N	3,163	3,346	1,200

Figures

Figure 1
UK Marginal Income Tax Rates

The figure plots the marginal personal income tax rate for individuals in employment and self-employment during the year 2007.

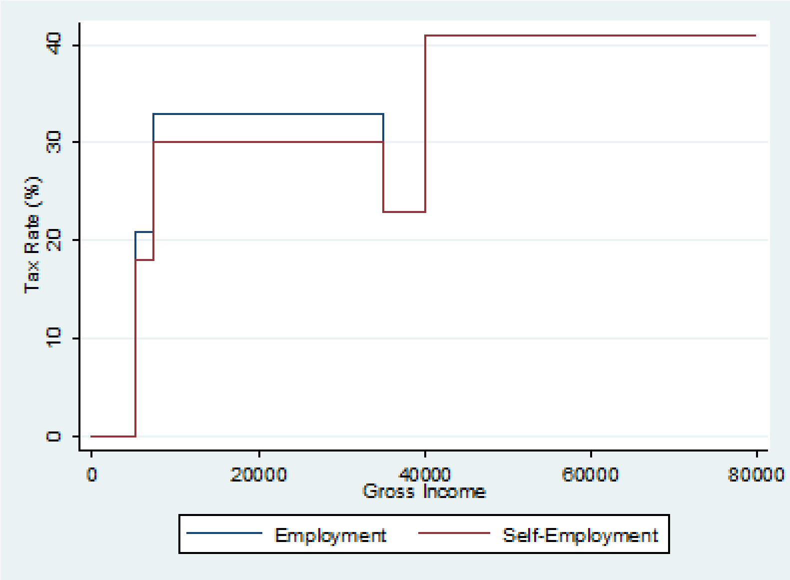

Figure 2
Convexity of Marginal Income Tax Rates

The figure plots total convexity, upside convexity and downside convexity for all incomes in the personal income tax code within the range £0 to £80,000 for the year 2007.

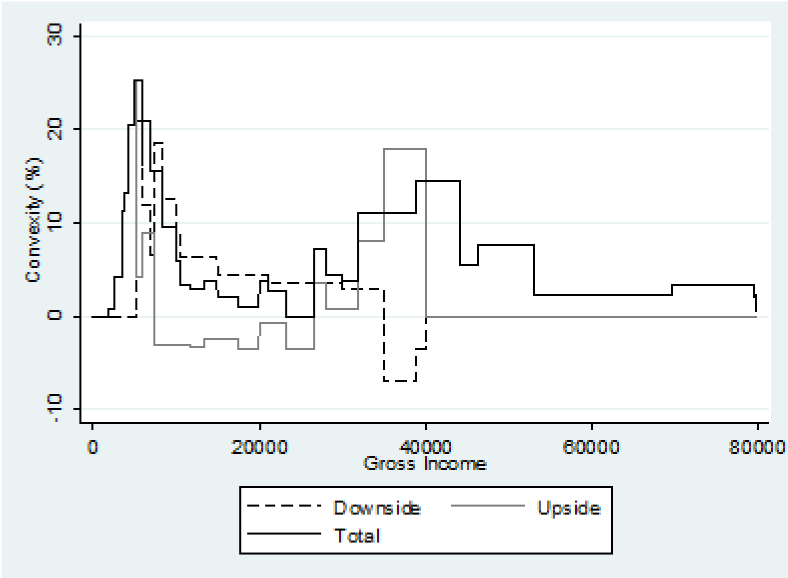

Figure 3
Tax Convexity Across Time

The figure plots total convexity of the all incomes in the personal income tax code within the range £0 to £80,000 for the years 2001, 2004 and 2007.

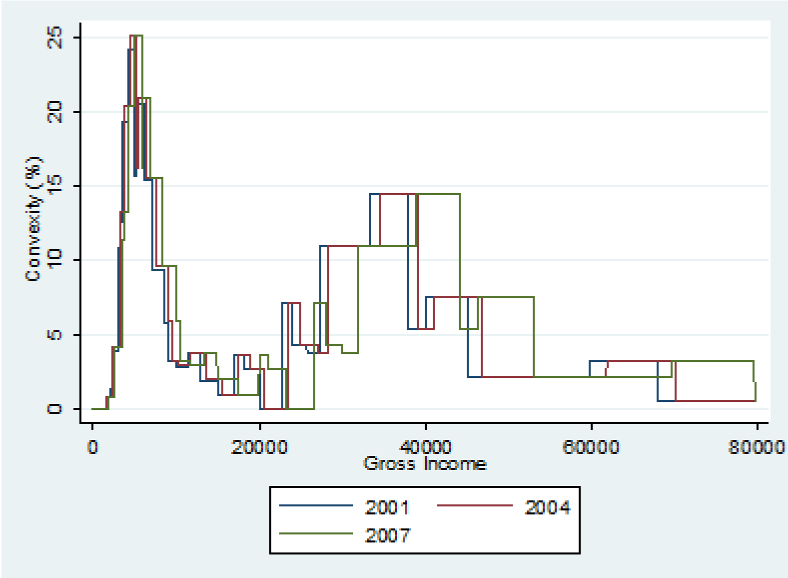

Figure 4
Upside and Downside Convexity Around the Top Tax Threshold (2007)

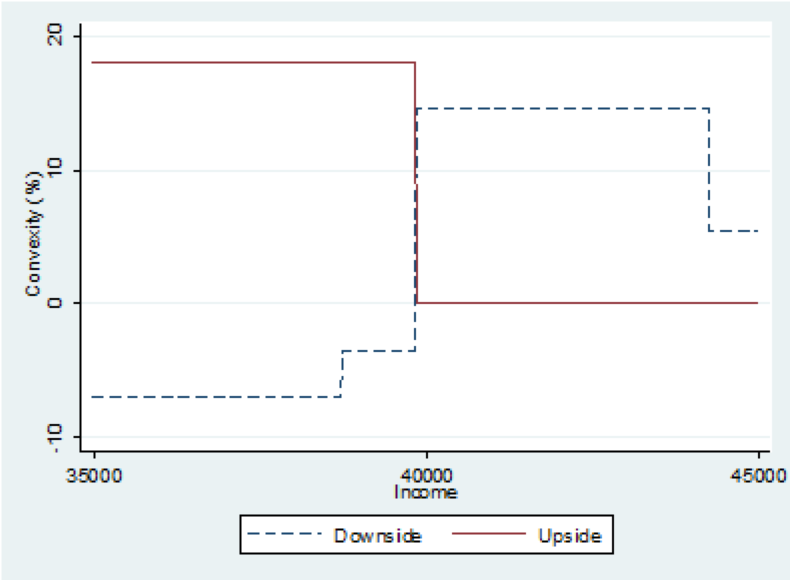

Figure 5
Average Tax Rate Paid by Employees (left hand axis) and Difference in the Average Tax Rate Paid by Employees and the Self-employed (right hand axis)

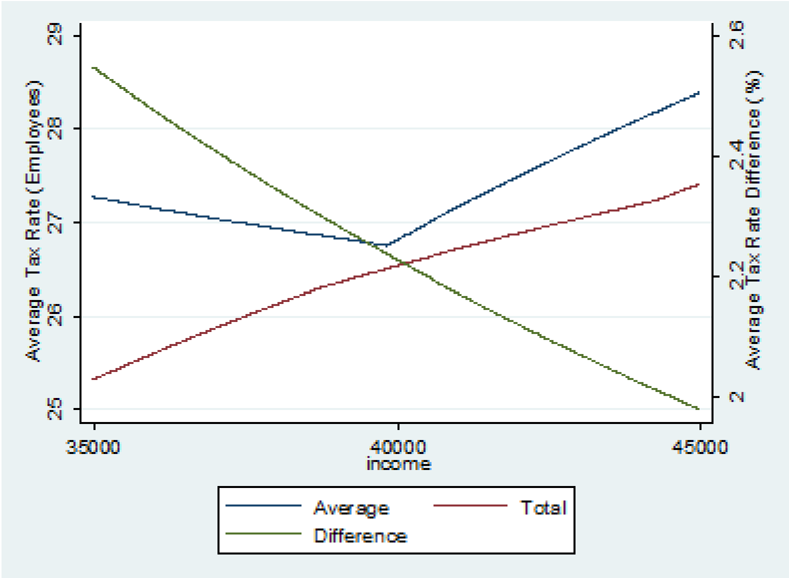

Figure 6A
Continuity of the Assignment Variable around the Cutoff

The figure plots a histogram and a kernel density function of the assignment variable within an income range of £10,000 either side of the HRT. The kernel density function is estimated using an epanechnikov kernel. The *x*-axis measures gross income minus the HRT ($x_{it} - x_0$). Data form all sample years are pooled.

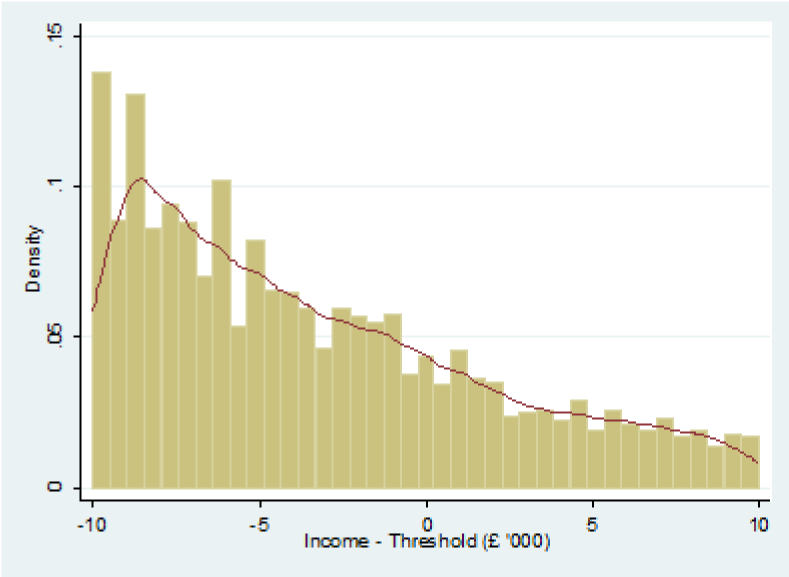

Figure 6B
McCrary (2008) Continuity Test

The figure plots kernel density functions using the method outlined in McCrary (2008) of the assignment variable over 1) the entire sample window within a £10,000 range either side of the HRT (solid black line); 2) the control group (black dashed line); 3) the treatment group (grey dashed line). In each case the kernel density functions are estimated using a triangle kernel. The *x*-axis measures gross income minus the HRT ($x_{it} - x_0$).

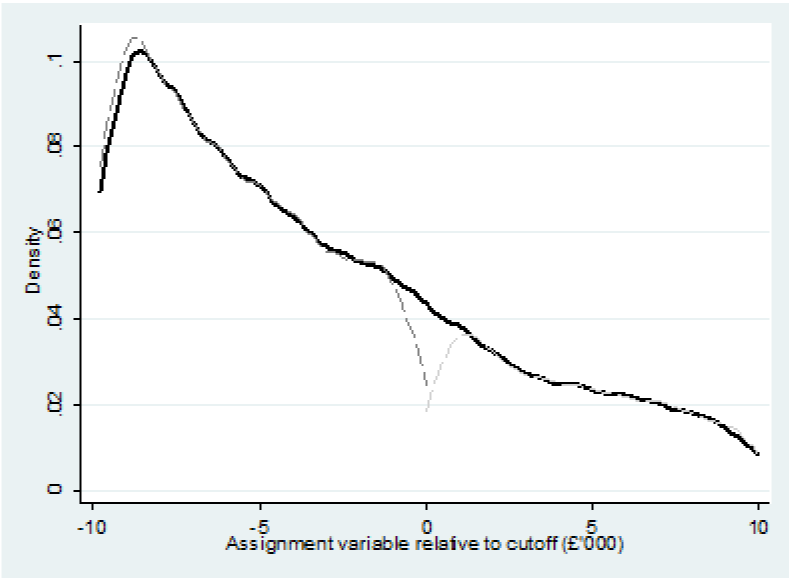

Figure 7
Entry into Self-Employment

The figure plots local polynomial regressions on the treatment and control group where the dependent variable is the mean rate of entry into self-employment within each £500-wide bin located within a £10,000 range either side of the HRT (*h* = £10,000). The local polynomial regressions are estimated using a triangle kernel. The *x*-axis measures gross income minus the HRT ($x_{it} - x_0$).

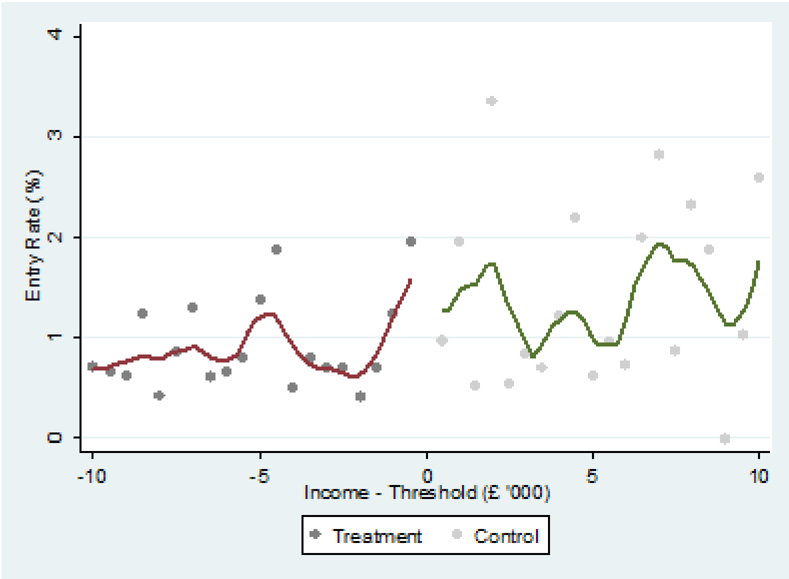

Appendix

Optimal Bin Width Calculation

To compute optimal bin width we follow convention in the literature by deciding to use K bins, corresponding to bin width h, based on visual inspection of the data (Lee and Lemieux, 2010). We then use a standard F-test that compares the fit of a regression model with K bin dummies to one where we further divide each bin into two equal sized smaller bins (i.e increase the number of bins to $2K$ and reduce the bin width to $h/2$). Because the first model is nested within the second, a standard F-test with K degrees of freedom can be used. If the null hypothesis is not rejected, this provides evidence that we are not oversmoothing the data by using K bins. Based on the evidence in the table below, the optimal bin width in our data is £500.

Appendix Table 1
Specification Test for Selecting Optimal Bin Width

The table presents the results of a test to select optimal bin width outlined by Lee and Lemieux (2010). Bin size is measured in £'s. R_r^2 (R_u^2) denotes the r-squared statistic from a regression using K ($2K$) bins. h = £10,000 in all rows. ** indicates statistical significance at the 1% level.

Bin size (£)	R_r^2	R_u^2	K	F-statistic	p-value
1,000	0.0012	0.0043	20	2.84**	0.01
500	0.0043	0.0087	40	1.07	0.40
250	0.0087	0.0138	80	0.60	0.99
125	0.0138	0.0271	160	0.79	0.95