

IZA DP No. 8606

**Big and Tall: Is there a Height Premium or
Obesity Penalty in the Labor Market?**

Wang-Sheng Lee

October 2014

Big and Tall: Is there a Height Premium or Obesity Penalty in the Labor Market?

Wang-Sheng Lee

*Deakin University
and IZA*

Discussion Paper No. 8606
October 2014

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Big and Tall: Is there a Height Premium or Obesity Penalty in the Labor Market?*

Previous studies have shown that both height and weight are associated with wages. However, by focusing on interpreting the partial effects of either height or weight on wages while holding all else constant, some gaps in our understanding of the complex relationship between body size and wages remain. Utilizing a semi-parametric spline approach, we first establish that a flexible analysis of height and weight provides a useful and meaningful proxy for beauty. A similar flexible analysis of height, weight and wages reveals that some combinations of anthropometric measurements attract higher wage premiums than others and that the optimal combination varies over the life cycle. A main contribution of the paper is in suggesting a novel and practical way of examining the returns to looks in the labor market based on objective anthropometric measurements.

JEL Classification: J31, J71

Keywords: height, P-spline, semi-parametric, wages, weight

Corresponding author:

Wang-Sheng Lee
Deakin University
School of Accounting, Economics and Finance
70 Elgar Road
Burwood, VIC 3125
Australia
E-mail: wang.lee@deakin.edu.au

* I am grateful to Fabrice Etilé, Dan Hamermesh, Christoph Wunder, Zhong Zhao, participants at the 5th Workshop on Economics of Health and Wellbeing in 2014 (Hepburn Springs, Australia) and participants at ESPE 2014 (Braga, Portugal) for helpful comments and suggestions. All errors are my own. This paper was previously presented at workshops using the title 'Physical Appearance and Wages over the Life Cycle.' This paper uses unit record data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The HILDA Project was initiated and is funded by the Australian Government Department of Social Services (DSS) and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views reported in this paper, however, are those of the author and should not be attributed to either DSS or the Melbourne Institute.

1. Introduction

There is a substantial literature in economics that focuses on the relationship between physical appearance and wages. One stream of this literature has focused on the relationship between height and wages (e.g., Loh, 1993; Persico et al., 2004; Case and Paxson, 2008; Case et al., 2009; Kortt and Leigh; 2010), generally finding a positive relationship between height and wages for both males and females. A second stream of this literature focuses on the relationship between obesity and wages (e.g., Register and Williams, 1990; Sargent and Blanchflower, 1994; Averett and Korenman, 1996; Pagan and Davila, 1997; Baum and Ford, 2004; Cawley, 2004; Han et al., 2009), generally finding a penalty for being overweight for women and mixed results for men.

More recently, in an attempt to improve on the literature that has been based on linear models researchers have also used semi-parametric methods to more flexibly examine the relationship between BMI and wages (e.g., Shimokawa, 2008; Hildebrand and Van Kerm, 2010; Gregory and Ruhm, 2011; Caliendo and Gerhsitz, 2014) and height and wages (e.g., Hübler, 2009; Hamermesh, 2012).

Curiously, both streams of this literature appear to exist in parallel universes. Most researchers choose to invoke the *ceteris paribus* assumption and focus on one aspect of body size at a time. This is likely because examining the relationship of a more complete measure of body size with wages requires the interaction of height and weight in a wage regression model to reflect various possible height and weight combinations. While the interpretation of interaction effects between two categorical variables or between one continuous variable and one categorical variable is rather straightforward, interaction terms between two continuous variables are generally avoided in practice by empirical researchers (e.g., see Lamina et al., 2012). Given this focus on interpreting the partial effects of height or weight on wages while holding all else constant, some gaps in our understanding of the relationship between body size and wages remain. For example, although we have empirical results suggesting that height is positively related to wages and that there exists a penalty for being overweight for women, the current literature is unable to provide immediate and clear answers to questions such as whether a tall and obese woman or a short and healthy weight woman would have a higher wage premium based on physical appearance, due to forces working in opposite directions.

The focus of this paper is on examining the non-parametric interaction effects of height and weight on wages. One might wonder how modelling this interaction in a flexible fashion improves on a standard OLS wage regression approach or the more recent semi-parametric approaches. Figure 1 summarizes representative findings from the literature superimposed on

a BMI chart, with height measured on the x-axis and weight on the y-axis. The figure is useful in illustrating what various approaches in the literature regard as key height and weight combinations that are associated with wage premiums.

In a standard linear regression model, suppose we are interested in comparing the wages of persons who are underweight, of healthy weight, overweight or obese and we use categorical variables to do so. Suppose further that we find a wage premium for overweight men, a typical empirical finding in the obesity and wages literature. In this case, the area between the iso-contour lines depicting BMI = 25 and BMI = 30 in Figure 1 will be the focal point of the figure as it represents all the possible height and weight combinations that are associated with a wage premium.¹ Note that this area includes both short and tall individuals.

[Figure 1 about here]

Using a semi-parametric approach, Gregory and Ruhm (2011) find using data from the Panel Study of Income Dynamics (PSID) that their conditional wage function peaks at BMI values of 22.8 for women and 26.7 for men. This is represented in Figure 1 by two thick BMI lines that are labelled “BMI = 23” and “BMI = 27”. The two lines highlight the various combinations of height and weight that correspond to the peak in expected wages for both genders. Similarly, Hübler (2009) reports that the maximum height effect is experienced by moderately tall men at 191 cm and moderately short women at 160 cm. His findings are represented by two vertical lines, one at 160 cm and another at 191 cm.

In a crude attempt to merge the results from these two separate literatures and focusing on the intersection of the lines in Figure 1, would it be fair to conclude that on average, a 1.60 meter tall woman who weighs 58.9 kg (BMI = 23) and a 1.91 meter tall man who weighs 98.5 kg (BMI = 27) are expected to have the highest wage premiums? Figure 1 is worth examining more closely as it highlights that a person’s physical appearance is a multi-dimensional construct and not necessarily a function of height and obesity considered separately. It will also be a very useful benchmark graphic for examining the results to be presented in this paper.

In the economics literature on physical appearance and wages, there is also a third stream that is based on using a more direct but subjective measure of physical appearance. For example, Hamermesh and Biddle (1994), Biddle and Hamermesh (1998) and Harper (2000)

¹ For example, the following heights and weights trace out the line corresponding to a BMI value of 25: (1.50m, 56.25 kg), (1.60m, 64 kg), (1.70m, 72.25 kg), (1.80m, 81 kg), (1.90m, 90.25 kg), (2.00m, 100 kg).

focus on examining the relationship between beauty and wages. In these studies, a beauty score is typically measured by asking interviewers in the field to rate the perceived attractiveness of survey respondents on a scale that varies from “very attractive” to “very unattractive.” Another approach is based on field experiments. In such studies, fictitious job applications with facial photographs attached are sent out to employers to examine the relationship between attractiveness and differential treatment in hiring in the labor market (e.g., Rooth, 2009; Boo et al., 2013). The magnitude of attractiveness effects is roughly the same as or larger than that of other important variables in the social sciences. For example, Hamermesh and Biddle (1994) find that the beauty premium (attractive vs unattractive) is equivalent to an extra 1.5 years of schooling.

It has been argued that BMI is a major factor in determining physical attractiveness for women (Tovée et al., 1998, 1999). Gregory and Ruhm (2011) also share the viewpoint that BMI is a proxy for physical attractiveness. They argue that the low levels of BMI at which their estimated wage function peaks for women makes it less likely that BMI measures the effects of obesity or poor health and more likely that it captures a beauty effect. Comparing directly the relative importance of BMI and the waist-to-hip ratio for women, Kościński (2013) finds that BMI is twice as important as the waist-to-hip ratio for the attractiveness of the female body.

Although BMI may be an important factor for female attractiveness, for men, Maisey et al. (1999) find that attractiveness is more determined by shape cues (specifically the waist-to-chest ratio, a measure of upper body shape) rather than BMI. Moreover, a single number such as the BMI might not convey enough information regarding body shape. For example, when we think of a particular BMI value for a person, it is difficult to visualize exactly what a person with a BMI value of 23 looks like given the vast number of possible height and weight combinations it allows. On the other hand, it is arguably easier to visualize someone who is 1.6 meters tall who weighs 59 kg (BMI = 23) or a person who is 1.7 meters tall and who weighs 66.5 kg (BMI = 23).²

The question we address in this paper is whether there is a more effective way of examining the relationship between physical appearance and wages using information on height and weight as a proxy for physical appearance.³ Subjective data on beauty is not

² This rationale explains why many internet dating sites such as match.com prefer to report information on height and weight separately rather than a person’s BMI.

³ Generally, joining a fitness club or gym to get into shape (especially during the new year) is more popular than resorting to plastic surgery to improve one’s physical appearance. Lee and Ryu (2012) find using Korean data that

routinely collected and available in most survey data sets unlike information on height and weight. The innovation in this paper is an attempt to analyse joint variation in height and weight without imposing any restrictive functional forms, such as that imposed by the BMI formula (weight/height²). Using seven waves of panel data on working adults from Australia as an illustration, this paper highlights how labor market returns to attractiveness, as proxied by height and weight, differ by gender and how these returns vary over the life cycle.

2. Data and Methods

2.1 Model

The main approach we use in this paper is to use panel data to estimate the following log wage equation:

$$W_{it} = f(Z_{it}) + X_{it}\beta + \bar{X}_i\eta + \alpha_i + \varepsilon_{it} \quad (1)$$

where W_{it} is the logarithm of the weekly earnings of individual i over time t and $Z_{it} = (H_i, W_{it})$ is an individual's height H_i and weight W_{it} . We assume that a person's adult height is time invariant and that shrinkage due to old age is not an issue.⁴ X_{it} is a vector of individual characteristics, α_i is an individual specific unobserved effect and ε_{it} is a random disturbance term of mean zero. In order to exploit the panel nature of the data set we use for our empirical analysis, we allow for correlated random effects by including the means of time-varying covariates \bar{X}_i in the model (Mundlak, 1978).

In equation (1), the function $f(\cdot)$ is a continuous but unspecified function of height and weight that is estimated from the data. The parametric part of the model allows discrete or continuous covariates that are deemed to have a linear effect on the outcome to be modeled alongside non-parametric terms. Recently, in estimating the effect of weight on wages, Kan and Lee (2012) allow for a more flexible semi-parametric specification of height and weight by adopting a semi-linear model consisting of smooth functions of weight and height and a linear function of the other regressors. They argue that such a specification provides deeper insights into the labor market returns to obesity relative to the conventional approach of

the associated monetary benefits from beauty premiums in labor and marriage markets are not high enough for most men and women to recover the surgery costs.

⁴ This rules out any fixed effects models since we are interested on the extent to which height matters in the labor market.

estimating a regression of log wages on BMI. By examining the weight effect on wages for white women in the NLSY at median height, they find that there is no weight effect on wage up to the average weight, beyond which a large negative effect kicks in. One key difference between their paper and this paper is that we place an emphasis on interpreting the interaction effects of height and weight, as compared to the effects of weight given height.⁵

The vector X_{it} includes a rich set of control variables. First, it includes variables for various socio-demographic characteristics. These include age (and its square), education, marital status, number of children (total number and number aged 0 to 4), region, state, country of birth (with two dummy variables included identifying whether the respondents is born overseas with English as their first language, or born overseas with English not their first language), and whether the respondent is an Indigenous Australian. Human capital and other job-related characteristics are also included. These include the variables years of experience in the current occupation (and its square), occupation (using 1-digit Australian and New Zealand Standard Industrial Classification 2006 codes), contractual employment status (with two dummies identifying casual and fixed-term contract employment, respectively), firm size (with dummies included to identify firms with between 1-99 employees and firms with between 100-499 employees), a quadratic in usual hours worked per week, and union membership. As a person's health is related to productivity, we also include the presence of a long-term health condition differentiated by whether it is work limiting or not. As a measure of family background and a proxy for the individual's socio-economic background, we also include the occupational status of the respondent's father at the time when he or she was aged 14.⁶ Finally, we also include a set of wave dummies.

In the existing literature on body size and wages, as there may be unobserved factors that affect wages that are correlated with body size (e.g., one's discount rate, self-discipline or social networks), emphasis has been placed on addressing the endogeneity of height and weight. While the endogeneity of weight and possible solutions to remedy it is an established issue in the literature (e.g., see the discussion in Cawley, 2004), a concern with the endogeneity of height is more recent. For example, some recent progress has been made on establishing

⁵ In non-linear models, as 'partial mean' or 'partial median' plots may be unattractive, Henderson et al. (2012) propose an alternative approach. They suggest the use of a '45° plot' where the advantage is that the plot no longer requires selection of where to fix the remaining variables when attention is placed on a specific variable.

⁶ The father's occupational status scale ranges from 0 (lowest-status occupations) to 100 (highest-status occupations). The scale takes account of the average education and income levels in an occupation (for more detail on the methodology, see Ganzeboom et al., 1992; Jones and McMillan, 2001).

causal links using genetic markers as instrumental variables to estimate causal effects of height (von Hinke Kessler Scholder et al., 2013).

Three approaches of dealing with the endogeneity of weight have been proposed in the literature. The first approach is to use lagged measures of weight to minimize the possibility of reverse causality (e.g., Averett and Korenman, 1996). The second is to use a fixed effects approach to eliminate time-invariant factors that affect both weight and wages (e.g., Baum and Ford, 2004). The third approach is to use the method of instrumental variables based on instruments such as a measure of a sibling's weight (e.g., Cawley, 2004). The focus in this paper is more on the flexible estimation of the relationship between height, weight and wages and less on the endogeneity of body size. The issue of endogeneity is addressed partially by two approaches used in this paper. The first approach involves the use of a correlated random effects model in equation (1), which is the main approach used in this paper. By including \bar{X}_i in the model and allowing for a correlation between α_i and X_{it} , Wooldridge (2010, chapter 10) shows that one obtains the fixed effects estimate in the context of a linear model. To the extent that modelling of the unobserved heterogeneity in this fashion helps capture any omitted variables, bias due to the endogeneity of body size will be reduced. The second approach is based on using lagged values of height and weight to mitigate potential reverse causality (see Section 4.5).

In the economics literature, a cross-sectional version of equation (1) which does not include the α_i and \bar{X}_i terms is often referred to as a partially linear regression model or a semi-parametric model because part of the model contains a parametric functional form but another part does not make any parametric assumptions. As highlighted by Ichimura and Todd (2007: 5419), such models have been broadly applied in economics but mainly to the problem of estimating Engel curves and the problem of controlling for sample selection bias. A widely cited semi-parametric model is the model proposed by Robinson (1988) which in its formulation treats the variables that enter the non-parametric part of the model $f(\cdot)$ as nuisance variables. Instead, Robinson's (1988) focus is on obtaining a \sqrt{n} -consistent estimator of the parameter vector β . Such an approach is very useful if one is interested solely in the parameter vector β but less useful if one is more interested in examining the effects of the variables that enter $f(\cdot)$.⁷

⁷ Chen (2007) refers to a model as being 'semi-parametric' if its parameters of interests are in finite-dimensional spaces but its nuisance parameters are in infinite-dimensional spaces, and a model as 'semi-nonparametric' if it contains both finite-dimensional and infinite-dimensional unknown parameters of interest.

As the model is additively separable and includes a non-parametric component, in the statistical literature, it is sometimes also referred to as a generalized additive model (GAM). This is because it extends a generalized linear model by replacing the linear functional form with an unknown functional form determined by the data. Lowess and smoothing splines can be very useful when there are no more than two explanatory variables. When there are more than two explanatory variables, GAMs provide a compromise between the ease of interpretation of the linear model and the flexibility of the general non-parametric model. Complicated non-linear problems can be easily accommodated, even for models with many explanatory variables. GAMs are able to accommodate the interaction of two or more predictors in a way that is conceptually comparable to interactions in a linear regression model. The joint smooth function of the predictors can be specified using tensor product smooths which is optimal for variables measured on different scales (Wood, 2006a). We use this interaction to map out the height-weight combinations that are associated with the highest weekly wages.

2.2 Data

The primary data used in this paper comes from the 2006-2012 waves of the Household, Income and Labour Dynamics in Australia (HILDA) survey, which are the waves where self-reported data on height and weight were collected.⁸ In a supplementary analysis discussed in Section 3 of the paper, we also use data from the National Longitudinal Study of Adolescent Health (Add Health) to make the case that a flexible analysis of height and weight provides a useful and meaningful proxy for physical appearance.⁹

HILDA is one of the largest surveys in Australia that is representative of the population and has detailed information on wages, body size and other relevant demographic characteristics such as education and marital status that are helpful in estimating labor market returns. The dependent variable we focus on is log weekly gross wages from all jobs for a person measured in constant 2011 dollars. Although it is possible to construct hourly wages

⁸ For a detailed discussion of the quality of the height and weight data in HILDA, see Wooden et al. (2008). Unfortunately, data on physical attractiveness (rated by the interviewer) is not available in HILDA. While physical attractiveness is largely a biological phenomenon, it can also be influenced by presentation such as the type of clothing, hair color and makeup (e.g., Hamermesh, Meng and Zhang, 2002).

⁹ Add Health is a program project designed by J. Richard Udry, Peter S. Bearman, and Kathleen Mullan Harris and funded by grant P01-HD31921 from the Eunice Kennedy Shriver National Institute of Child Health and Human Development, with cooperative funding from seventeen other agencies. Special acknowledgment is due Ronald R. Rindfuss and Barbara Entwisle for assistance in the original design. Persons interested in obtaining data files from Add Health should contact Add Health, Carolina Population Center, 123 W. Franklin Street, Chapel Hill, NC 27516-2524 (addhealth@unc.edu).

from the data, the reason we focus on weekly wages is to minimize measurement error in the outcome. As hours of work is a variable potentially measured with error, the use of weekly wages helps avoid the issue of division bias (Borjas, 1980). Furthermore, in each wave of the HILDA survey, approximately 40% of individuals report their wages based on the previous week and another 40% report it based on the most recent fortnight. This reflects the system of wage payments in Australia and the use of weekly wages is therefore closest to its original reported form in the data. In order to correct for potential reporting error in height (i.e., where there are discrepancies across waves by person), we correct the raw data manually for obvious inconsistencies and use the average of a person's height across waves in our regression models.

We restrict our sample to respondents aged 25–54 who are employed and work at least 20 hours per week in their main job, dropping those respondents who are self employed, enrolled in full-time education and who have missing data on individual characteristics used in our model.¹⁰ Persons with implausibly low weekly wages (i.e., less than \$10, which is significantly below the federal minimum wage) are dropped. To minimize instability of the statistical model at extreme height and weight values, height and weight were trimmed at 0.1% and 99.9% of their respective gender distributions. Males < 1.55 meters or > 2.00 meters, as well as weighing < 50 kg or > 173 kg were dropped from the analysis, whereas females who were < 1.44 meters and > 1.86 meters and weighing < 40 kg or > 160 kg were dropped from the analysis. The unbalanced panel data used for our empirical analysis comprises of 13,803 male observations over time based on 4,060 males, and 11,944 female observations over time based on 3,845 females.

Table 1 provides descriptive statistics of the pooled data from waves 6 to 12 of HILDA used in this paper. Since previous studies have sometimes observed differences in the effect of body size on the wages of men and women, we estimate our results separately for men and women. Weekly earnings in constant 2011 dollars is \$1508 for males in our sample and \$1071 for females. The average height is 1.79 meters for men and 1.65 meters for women, while the average weight is 87.0 kg for men and 71.8 kg for women. This makes both the average man and woman in Australia slightly overweight according to BMI.

[Table 1 about here]

¹⁰ We follow a large proportion of the literature on obesity and wages in not focusing on the issue of selection into employment, implying that the results we obtain are not generalizable to the overall population.

2.3 P-Splines

In this paper, we focus on using P-splines for performing our empirical analysis (Eilers and Marx, 1996). P-splines are a simple combination of two ideas for curve fitting: the use of B-splines as the basis function with the inclusion of a difference penalty on the regression coefficients. Marx and Eilers (1998) further introduce P-splines to the GAM setting. As an alternative, the function $f(\cdot)$ in equation (1) can be constructed using kernel functions that require marginal integration (e.g., Fan and Gijbels, 1996; Li and Racine, 2007, Kan and Lee, 2012). We do not pursue this alternative approach in the paper.

The asymptotic behavior of P-spline estimation has been explored in detail recently. Hall and Opsomer (2005) use a white-noise process representation of P-splines to provide insights into its asymptotic properties. Li and Ruppert (2008) provide theoretical asymptotic results where they derive equivalence between kernel smoothing and penalized splines in the univariate case with a large number of knots. Claeskens et al. (2009) relate the asymptotic properties of P-splines to known asymptotic results for regression splines and smoothing splines. Relaxing the simplifying assumption that the dimension of the spline basis is fixed which characterized earlier theoretical work, Kauermann et al. (2009) show that P-spline smoothing is asymptotically justified even if the number of spline basis functions is allowed to increase with the sample size. We briefly describe the two components of P-splines below as they are likely to be unfamiliar to economists.

Under a B-spline approach, the unknown $f(\cdot)$ function can be written as the sum of k basis functions:

$$f(x) = \sum_{i=1}^k \alpha_j b_j(x) \quad (2)$$

For expositional purposes, an example of a basis function that has flexible turning points is the polynomial basis:

$$f(x) = \alpha_1 + \alpha_2 x + \alpha_3 x^2 + \alpha_4 x^3 + \alpha_5 x^4 \quad (3)$$

Such a function will clearly fit data better than simply using linear or quadratic terms in a linear regression model. For example, Murphy and Welch (1990) argue that allowing for quartic terms in experience is empirically important and superior than using a quadratic in experience for fitting the earnings curve. Although the polynomial basis is useful for illustrating ideas, in

practice, it has poor numerical stability properties as they tend to lead to highly correlated parameter estimators. Instead, B-splines are polynomial pieces that are joined together in a special way. A B-spline function is a piecewise polynomial function of degree k and the places where the pieces meet are referred to as knots. As B-splines are continuous functions at the knots, it is possible to control the wiggleness of the fitted model by controlling the number of piecewise polynomial functions used.¹¹ As discussed by Eilers and Marx (1996), when a large number of equally spaced knots and a large number of B-splines are used, the primary role of the basis function is to serve as a convenient smooth interpolation device.¹²

P-spline smoothing models are fit using penalized likelihood maximization in which the model likelihood is modified by the addition of a penalty for each smooth function, penalizing its ‘wiggleness’. Specifically, Eilers and Marx (1996) propose adding the following penalty to the objective function to be minimized when using B-splines:

$$P_{\lambda}(f) = \lambda \sum_{i=d+1}^k (\Delta^d \beta_i)^2 \quad (4)$$

The penalty can be written as a linear combination of some basis functions, where Δ^d is the difference operator of order d .¹³ To control for the tradeoff between penalizing wiggleness and penalizing badness of fit, each penalty is multiplied by an associated smoothing parameter. P-splines are not troubled by boundary effects, which is an issue that plagues many types of kernel smoothers. It is, however, important to specify the domain of the data correctly prior to estimation (Eilers and Marx, 1996: 98).

In the economics literature, Wunder et al. (2013) is a recent application that focuses on examining life satisfaction as a smooth function of age over the life-span. Clark and Etilé (2011) use P-splines to examine how the interaction of own BMI and partner’s BMI affect one’s life satisfaction. Wu and Sickles (2013) perform Monte Carlo simulations and report on the finite sample performance of P-spline methods. They also illustrate its usefulness in capturing shape restrictions in economics applications, which help provide relative flexibility in fitting non-linear relationships.

¹¹ A good overview of B-splines is provided in De Boor (2001).

¹² A competing approach to smoothing described in Ruppert and Carroll (2000) and Ruppert and Wand (2003) is based on truncated power functions.

¹³ An alternative approach to using P-splines is based on thin plate regression splines (Wood, 2003) which are low rank smoothers which avoid having to choose knot locations for spline bases. However, this approach has been shown to perform less well in Monte Carlo simulations compared to the P-spline approach we use in this paper (Strasak et al., 2011).

One obvious drawback of spline smoothers such as the B-spline smoother is that the number and position of the knots has to be chosen by the researcher. If too many knots are used such that the smoothing spline can approximate the unknown true function, it is likely that the model will overfit data that contains any noise. Eilers and Marx (1996) recommend using a large number of equally spaced knots (e.g., between 10 and 30). Ruppert (2002) notes that the choice of the number of knots is usually not critical in affecting the model results as long as it is set large enough. This is because it only sets an upper bound on the flexibility of a term. Instead, the smoothing parameter plays a more important role in controlling the actual effective degrees of freedom.

As is the case with many semi- and non-parametric approaches, the choice of the smoothing parameter λ or the amount of smoothing that is applied to the data strongly affects the fit of the model. The smoothness of $f(\cdot)$ is calculated with the aim of optimal balance between the fit to the data versus a penalty for excessive “wiggleness” of the functions. Several approaches to choosing λ are available. One possibility is to estimate it based on some goodness of fit criterion, such as AIC (Eilers and Marx, 1996). Another possibility is to use a data driven approach such as generalized cross validation (Wahba, 1990). Finally, the smoothing parameter can also be estimated via a restricted maximum likelihood (REML) approach (Ruppert et al., 2003; Wand, 2003).

In this paper, we estimate the smoothing parameter using REML. In the presence of correlated data (which applies to our context which is based on repeated observations on individuals), standard smoothing parameter selectors fail to work and tend to overfit the data (Opsomer et al., 2001). Krivobokova and Kauermann (2007) show in simulations and a theoretical investigation that REML-based smoothing parameter selection is less sensitive to misspecifications of the correlation structure in the data compared to approaches that minimize the mean-squared error, such as generalized cross validation and those based on the AIC. Furthermore, the Monte Carlo simulation results in Strasak et al. (2011) suggest that compared to the alternative approach of using generalized cross validation, REML for smoothing parameter selection in conjunction with the use of P-splines seems to provide the best compromise between goodness of fit and stability of the estimator.

In equation (1), rather than estimating the model based on pooled data, we exploit the panel structure of the data and estimate a generalized additive mixed model (GAMM), which allows the inclusion of individual specific random effects. In our application, we use cubic B-splines as the basis with a second order difference penalty. We use a basis dimension of 5 for

both height and weight, for a total dimension of 25. The mgcv library (Wood, 2006b) in R version 3.0.2 is used to estimate the models.

3. Beauty Matters

Drawing widely from the anthropological, psychological, biological and archeological literature, Etcoff (2000) discusses how beauty is an essential and ineradicable part of human nature that is revered and ferociously pursued in nearly every civilization. Dion et al. (1972) in the groundbreaking study titled “What is Beautiful is Good” investigated the perceptions of physical appearance within society. Their paper concluded that a physical attractiveness stereotype was observed as statistically significant and that this identified stereotype has the potential to impact a large variety of social interactions and to have a substantial influence on the behavior of individuals. In a meta-analysis, Langlois et al. (2000) present a summary of studies examining the “What is Beautiful is Good” phenomenon and report that physically attractive individuals are more likely than unattractive individuals to be judged as competent in their professions, to experience success in their occupations, and to be treated more favorably by others.

3.1 Depicting beauty using height and weight combinations

Evolutionary psychologists have long argued for the existence of universally shared criteria of attractiveness and beauty, despite the fact that beauty is largely a subjective phenomenon. In this paper, we attempt to proxy for physical appearance using information on both height and weight in an interactive fashion. This is a natural extension of the relationship found between BMI and attractiveness found for women (Tovée et al., 1998, 1999) without being restricted by the functional form imposed by the BMI formula. In order to demonstrate the suitability of using information on the interaction of height and weight to proxy for physical appearance, we use data from Wave 4 of the National Longitudinal Study of Adolescent Health (Add Health) to examine the relationship between height, weight and an interviewer rated score assessing the beauty of the respondent. In this analysis, we illustrate how subjective beauty scores are related to various height and weight combinations, and why the latter can be viewed as a practical alternative measure.

Interviewers for Add Health were asked the following question: “How physically attractive is the respondent?” Possible answers were (1) very unattractive, (2) unattractive, (3) about average, (4) attractive, and (5) very attractive. We create a binary variable where beauty = 1 if the respondent has a rating of 4 or 5 and beauty = 0 if the respondent has a rating of 1, 2

or 3. The Wave 4 public-use file for Ad Health contains data on 5,114 respondents, aged 24 to 32.

In order to estimate the effects of height and weight on beauty, we estimate a variation of equation (1) where the focus is on modelling beauty as a binary outcome variable.

$$Attractiveness_i = f(Z_i) + X_i\beta + \varepsilon_i \quad (5)$$

In this generalized additive model, we assume a logistic link function and that the distribution of the outcome is binomial. Covariates included in this model based on the Add Health data are age, education, race, household income and employment status.

In general, the non-linear results depicting the unknown function $f(\cdot)$ are best shown using graphs as the complex non-linear interaction effect between height and weight cannot be simply summarized using a regression coefficient or mathematical function. The joint interactive effect of height and weight can be shown in a three-dimensional plot or a contour plot as they both contain the same information. We choose to use the latter for our application because the iso-contour lines in the plots make it very easy to see how the probability of being classified as beautiful is associated with different combinations of height and weight.

Figure 2 reveals that height is a dominant factor in determining physical attractiveness of young males with taller and proportionately heavier males deemed to be the most attractive. On the other hand, being slim matters more for young females. Figure 3 suggests that young females who are tall and slim (but not too slim and below 50 kg) are considered to be the most attractive.¹⁴

[Figure 2 about here]

[Figure 3 about here]

4. Results

4.1 Linear parametric models

In estimating the relationship between height/weight on labor market outcomes in a multivariate setting, there is some disagreement in the literature over whether one should also control for education, experience and marital status. The central issue is whether experience

¹⁴ Tovée et al. (1998) had found in a regression of attractiveness ratings on a cubic polynomial in BMI that females with a BMI value of 20 had the highest predicted attractiveness ratings. To illustrate the extra information the function $f(\cdot)$ conveys in a contour plot, Figure 3 should be compared with the BMI = 20 contour line in Figure 1.

and education are a channel through which body size affects earnings, or whether they are confounding variables that happen to be correlated with both body size and wages. Some authors argue against accounting for differences in decision variables when estimating the effect of labor market discrimination (e.g., Neal and Johnson, 1996; Heckman, 1998). This is because including variables such as education, experience and marital status which are subject to worker choice and which can be affected by labor market discrimination will lead to the wage effects of discrimination being misstated. If the goal is to estimate the total effect of height/weight on wages and not the partial effect of height/weight on wage conditional on education, experience and marital choices, then one might prefer to omit these variables from the model, controlling only for purely exogenous determinants of wages. However, other authors deem it might also be of interest how much these factors can account for the variation in wages (e.g., Cawley, 2004; Baum and Ford, 2004). In this paper, although we estimated results for both sets of models, we focus on presenting the results of models where educational attainment, work experience and marital status are included.¹⁵

For benchmark purposes, we first estimate linear models to illustrate the relationship between body size and wages using our data. Table 2 provides results for males. Columns 1 to 4 are based on using pooled data and OLS. Columns 5 to 8 are based on using a correlated random effects approach and can be viewed as the linear version of equation (1).¹⁶ The results in Table 2 are generally consistent with what is found in the literature. There is a height premium for males (column 1), with a 10 cm increase in height corresponding to a 3.4% increase in wages. Weight controlling for height (column 2) and BMI (column 3) do not appear to be significantly related to wages. While at odds with the results for most other countries, the insignificance of the BMI coefficient here is consistent with Kortt and Leigh's (2010) findings using the 2006 and 2007 waves of the HILDA data.

When clinical weight classifications based on BMI cutoffs are used, overweight men appear to have a 3.3% wage premium over healthy weight men holding all else constant. Results based on the correlated random effects model are generally similar, with the exception that the results for overweight men based on BMI cutoffs are no longer statistically significant.

[Table 2 about here]

¹⁵ The estimation results where education, experience and marital status are excluded are very similar and available upon request.

¹⁶ Note that as height is not time-varying in our data set (we use a person's average height measured across all waves), it is not possible to estimate a linear fixed effects model when the coefficient on height is of interest.

Table 3 provides the OLS results for females. In this case, while there is still a positive return to height (column 1), there is now a 2.0% wage penalty associated with obesity (column 4). Both these findings are again consistent with the general literature. However, the coefficient on obesity loses its significance in the correlated random effects model.

[Table 3 about here]

4.2 *Unconditional smooths of body size and wages*

The use of linear models to examine the relationship between body size and wages might not be ideal if there exist non-linearities in the relationship. This has been the motivation behind researchers using semi-parametric approaches in examining the relationship between BMI and wages (e.g., Shimokawa, 2008; Hildebrand and Van Kerm, 2010; Gregory and Ruhm, 2011; Caliendo and Gehrsitz, 2014). Figure 4 illustrates the unconditional non-linear relationship between body size and log hourly wages in the HILDA data for men and women.¹⁷ In the top panel, the relationship between height and wages is examined. Our results using Australian data is similar to what is reported in Hübler (2009) using German data. There is evidence that a linear function might not be appropriate for both men and women. The bottom panel examines the bivariate relationship between BMI and wages. For men, wages appear to peak at a BMI of between 25 and 32 whereas for women, wages start to dip noticeably after a BMI value of 22. These findings generally echo the results found in this literature.

[Figure 4 about here]

4.3 *General additive models – full sample*

Based on using equation (1), Figure 5 depicts the height and weight combinations of males aged 25-54 that are associated with the highest log weekly wages. In general, the iso-contour lines suggest the existence of a height premium over a broad weight range. For example, in Figure 5, looking at males who weigh 100 kg, it can be seen that moving to the right in the figure (which represents an increase in height holding weight constant) leads to higher log weekly wages.

¹⁷ We use log hourly earnings here to allow comparability with much of the literature that focuses on log hourly wages as the outcome.

In comparison with a semi-parametric analysis that examines the non-linear relationship between BMI and wages, we argue that these results present a more complete and realistic picture. For instance, suppose we focus on a result from Gregory and Ruhm (2011) that suggests that males wages peak at a BMI value of 27. If we superimpose a BMI = 27 line in Figure 5 (i.e., compare Figures 1 and 5 side by side), it can be seen that the BMI = 27 line is part of the story but not the whole story. Instead, it appears that there is a wider combination of height and weight that is associated with the highest levels of wages, as depicted by the highest iso-contour line. Weight also does matter but it appears that it matters only to the degree which it is deemed to be proportional relative to height. For example, a male who is 1.80 meters tall maximizes his wages if he weighs between 70 kg to 100 kg and a male who is 1.95 meters tall maximizes his wages if he weighs between 95 kg to 110 kg. These weight ranges correspond to the boundaries defined by the iso-contour lines when we move vertically up the figure, holding height constant.¹⁸

[Figure 5 about here]

The results for females shown in Figure 6 also suggest that increases in height are highly associated with higher wages. However, there is a wider range of heights and weights where expected wages are close to its highest levels. It is interesting to note that part of Figure 6 also encompasses the results found in Kan and Lee (2012) who find that for white women of average height in the NLSY, there is no weight effect on wage up to the average weight. Drawing a vertical line up from 1.65 meters (the average height of women in our sample) in Figure 6, we can see that wages stay constant but start decreasing when weight reaches close to the average weight of women in the sample.

[Figure 6 about here]

¹⁸ In creating a three-dimensional plane depicting the interaction between height, weight and log wages, there is a concern with the possibility of extrapolation and results at certain combinations of height and weight not supported by the data. This implies that the contour plots might not be relevant to any actual cases of the observed data. The scatterplot in Figure A.1 in the appendix highlights the various combinations of height and weight that exist in the sample, with each point representing an individual. The results in the contour plots presented in this paper are restricted to areas where there is thick support in the data, as one can readily ascertain by comparing the respective contour plot with the corresponding scatterplot in Figure A.1.

4.4 General additive models by age subgroups

The importance of a person's looks likely matters more at certain points in a person's life than others.¹⁹ Individuals also tend to put on weight as they age. As the contour plots presented in Figures 5 and 6 cover individuals over a broad age range (25-54), they potentially mask interesting heterogeneous effects that occur at different life stages. Hence, we also apply both these model specifications for different age groups in the data because we are interested to know how the importance of physical attractiveness and its link with wages varies over the life course. The motivation behind this is that a person's physical appearance might be more important at the earlier stages of one's life where careers are being built, but less important at the later stages of one's life where careers have stabilized and discrimination on looks matters less. Early in one's career, workers are to some extent considered to be unknown quantities to their employers, and physical appearance might therefore play a larger role in assessing their productivity. In addition, more physically attractive persons might also gain access to valuable networks that will help them climb the career ladder more quickly.

For males aged 25-39 (Figure 7), it can be seen that a height premium exists and that taller men generally earn the highest wages. These wage differences are economically significant. For example, consider a male aged 25-39 in the top panel of Figure 4 who is 1.90 meters tall and weighs 100 kg. Compared to a shorter male who is 1.70 meters tall but who is of the same weight, the taller person earns a wage premium of 0.1 log points or approximately 10%. According to the OLS model whose key results were presented in Table 2, this is almost economically equivalent to a difference between having a high school degree versus not having a high school degree.

[Figure 7 about here]

For middle-aged males who are aged between 40-54 (Figure 8), the existence of a height premium seen for young men persists. This is consistent with Figure 5, where it was shown that for the full sample of males aged between 25-54, a clear height premium exists. However, the figure also suggests that having a proportionate weight for height is not so important anymore and that many tall and heavy men are doing well in the labor market, consistent with a 'portly banker' hypothesis.

¹⁹ For instance, physical appearance will be relatively important for young people seeking to get married and for divorcees looking for new romantic partners.

[Figure 8 about here]

The appearance-wage dynamics for women over the life cycle differs from what we have seen for males. For young women aged 25-39 (Figure 9), the slopes of the iso-contour lines suggest that there are clear returns to being tall and slim up to a height of about 1.70 meters (which is relevant for about 80% of women in the sample who are less than 1.70 meters tall). Both height and weight appear to be equally important as can be inferred from the iso-contour lines which are close to having a 45 degree slope in some instances. Again, by comparing these results with Figure 1 side by side, one can see that the line of height and weight combinations giving rise to a BMI value of 23 in Figure 1 tells a different story.

[Figure 9 about here]

By middle age (40-54), the importance of having a particular body type seems to have dissipated somewhat. There are some women who are taller and heavier that do well in the labor market (see top right of Figure 10). There is also some evidence that being slim (weighing less than 80 kg) is important over a wide range of values for height. We speculate that one possible explanation for this finding in Figure 10 is that as women aged 40-54 likely face pressures from a combination of both age and gender discrimination, those that make a special effort in their appearance tend to be more successful in the labor market. For example, Lahey (2008) finds in a field experiment that younger women are more than 40% more likely to be offered an interview than women over 45. Supporting such a concern is the fact that Porter (2003) makes the case that there is a serious shortcoming in employment discrimination laws for older females. She argues that in order to protect older female workers, sex discrimination and gender discrimination should not be analysed separately but instead treated as a combined subclass in the law in the same way that sex and race is.²⁰

[Figure 10 about here]

²⁰ As a robustness check, we also tried augmenting the model with possible omitted variables such as personality and cognitive measures. Due to missing observations on these measures, the sample size was reduced considerably and for that reason we do not include them in our main specification. It is reassuring, however, to note that the results of the paper remain robust to the inclusion of these variables in this reduced sample.

4.5 Lagged height and weight

As past and current body shape and physical appearance can have an effect on current labor market outcomes, it might be of interest to examine how height and weight from several years ago affects one's wages in the current period. As in Averett and Korenman (1996), Gregory and Ruhm (2011) and Chen (2012), using such an approach allows one to mitigate any potential issues with reverse causality. Although current wages might possibly have an effect on current weight (e.g., low-wage workers who are depressed eat more junk food and become fatter), current wages are unlikely to have an effect on past weight.

In order to empirically examine how past body shape affects current wages in our data, we estimate a variation of equation (1) where we restrict the data to be a cross-section from the latest wave – wave 12 that corresponds to the year 2012. Instead of using height and weight measured in 2012, however, we use height and weight measured in 2006, the earliest wave where data on height and weight is available.

$$W_{i,2012} = f(Z_{i,2006}) + X_i\beta + \varepsilon_i \quad (6)$$

This provides us with an estimate of how body shapes have an effect on labor market outcomes six years later.

Based on the full model specification that includes education, experience and marital status, Figure 11 shows that men aged 25-54 in 2006 who are taller appear to have done the best six years later. On the other hand, although height also matters, the women in 2006 who earn the highest log weekly wages in 2012 are those who are tall and slim (Figure 12).²¹ Both results are consistent with the results we have seen for the age subgroups previously as well as the Add Health results for beauty.

[Figure 11 about here]

[Figure 12 about here]

²¹ There might be certain age ranges where physical appearance has the largest impact on labor market outcomes. For example, Chen (2012) finds that a higher BMI in one's twenties for females has a persistent negative effect on wages and is more important factor than current BMI. Unfortunately, the sample sizes in our are too small to reliably estimate such a model.

5. Conclusions

The literature on anthropometric measures and wages has been loosely linked to the literature on beauty and wages; researchers often speculate that results in the former provide some evidence of beauty or appearance effects. A main contribution of this paper is that it provides supporting evidence that links these two distinct streams in the literature. In utilizing a semi-parametric spline approach that allows for complex non-linear interactions, we suggest a novel and practical way of examining the beauty premium in the labor market based on objective anthropometric measurements.

The analysis of beauty as an outcome variable in the Add Health data suggest that a flexible analysis of height and weight provides a useful and meaningful proxy for physical appearance. A similar flexible analysis of height and weight on wages using nationally representative Australian data suggests that there are different returns in the labor market to various combinations of height and weight, and that the optimal combination of height and weight varies as one ages. Taken together, these results provide empirical evidence that wage premiums based on physical appearance exist and that the returns to looks change over the life cycle.

The approach employed in this paper helps provide a more complete understanding of how anthropometric measures matter in the conditional earnings function. Compared to results from a linear regression model, it provides a more complete picture of how body shapes that represent actual human forms are related to wages. Different strands of the literature on body size and wages have documented either an obesity penalty or a height premium. However, it is really a combination of height and weight that should be of interest. As weight tends to vary with height simultaneously, it can be problematic to interpret partial effects in a linear regression model when both height and weight or height and BMI are included in a wage equation. The current literature is not able to provide clear answers to questions such as whether there is a premium or penalty for a tall and obese person, or similarly whether there is a premium or penalty for a short and non-obese person. For example, although the literature suggests that women with relatively low BMI values tend to have the highest wages, one needs to keep in mind that such BMI values can correspond to women who are very tall or very short. The results in this paper make it clear that it is the combination of being tall and slim that is the driving force behind the beauty premium.

The emphasis of this paper has been on the use of a flexible econometric approach to estimate the effect of body size on wages. The estimates provided in this paper should not be regarded as the causal effects of height and weight on wages. Instead, they capture the residual

effect of height and weight on wages after controlling for relevant individual characteristics and can be regarded as a form of labor market discrimination against less physically attractive individuals. Appearance-based discrimination is an issue examined in more detail in Rhode (2010).²² It will be interesting in future work to attempt to search for evidence of appearance-based discrimination in particular sectors, as well as in different countries, using flexible estimation approaches such as the one used in this paper.

²² Writing from the legal perspective, Rhode (2010) provides a systematic survey of how appearance laws work in practice, and strategies for ending appearance-based discrimination.

References

- Averett, S. and S. Korenman. (1996). The economic reality of the beauty myth. *Journal of Human Resources*, 31: 304–330.
- Baum, C. and W. Ford. (2004). The wage effects of obesity: a longitudinal study. *Health Economics*, 13: 885–99.
- Boo, F., M. Rossi and S. Urzúa. (2013). The labor market return to an attractive face: Evidence from a field experiment. *Economics Letters*, 118: 170–172.
- Borjas, G. (1980). The relationship between wages and weekly hours of work: the role of division bias. *Journal of Human Resources*, 15: 409–423.
- Caliendo, M. and M. Gehrsitz. (2014). Obesity and the Labor Market: A Fresh Look at the Weight Penalty. IZA Discussion Paper No. 7947.
- Case, A. and C. Paxson. (2008). ‘Stature and status: height, ability, and labor market outcomes. *Journal of Political Economy*, 116: 499–532.
- Case, A., C. Paxson and M. Islam. (2009). Making sense of the labour market height premium: evidence from the British Household Panel Survey. *Economics Letters*, 102: 174–176.
- Cawley, J. (2004). The impact of obesity on wages. *Journal of Human Resources*, 39: 451–474.
- Cawley, J., K. Joyner and J. Sobal. (2006). Size matters: the influence of adolescent’s weight and height on dating and sex. *Rationality and Society*, 18: 67–94.
- Chen, X. (2007). Large sample sieve estimation of semi-nonparametric models. *Handbook of Econometrics*, Vol. 6B, pp. 5548-5632, Heckman, J. and Leamer, E. (eds.), North-Holland: Amsterdam.
- Chen, A. (2012). When does weight matter the most? *Journal of Health Economics*, 31: 285–295.
- Claeskens, G., T. Krivobokova and J. Opsomer. (2009). Asymptotic properties of penalized spline estimators. *Biometrika*, 96: 529–544.
- Clark, A. and F. Etilé. (2011). Happy house: spousal weight and individual well-being. *Journal of Health Economics*. 30: 1124–1136.
- De Boor, C. (2001). *A practical guide to splines*. Springer, New York.
- Dion, K., E. Berscheid and E. Walster. (1972). What is beautiful is good. *Journal of Personality and Social Psychology*, 24: 285–290.
- Eilers, P. and B. Marx. (1996). Flexible smoothing with B-splines and penalties (with comments and rejoinder). *Statistical Science*, 11: 89–121.

- Etcoff, N. (2000). *Survival of the prettiest: the science of beauty*. Anchor: New York.
- Fan, J., and I. Gijbels. (1996). *Local polynomial modelling and its applications*. London: Chapman and Hall.
- Ganzeboom, H., P. DeGraaf and D. Treiman. (1992). A Standard International Socio-Economic Index of Occupational Status. *Social Science Research*, 21: 1–56.
- Gregory, C. and C. Ruhm. (2011). Where does the labor market penalty bite? in *Economic Aspects of Obesity*, ed. by M. Grossman, and N. Mocan, chap. 11, 315–347, Chicago University Press, Chicago.
- Hall, P. and J. Opsomer. (2005). Theory for penalised spline regression. *Biometrika*, 92: 105–118.
- Hamermesh, D. (2012). Tall or taller, pretty or prettier: is discrimination absolute or relative? *IZA Journal of Labor Economics*, 1:2.
- Hamermesh, D. and J. Biddle. (1994). Beauty and the labor market. *American Economic Review* 84: 1174–1194.
- Hamermesh, D., X. Meng and J. Zhang. (2002). Dress for success – does priming pay? *Labour Economics*, 9: 361–373.
- Han, E., E., Norton and S. Stearns. (2009). Weight and wages: fat versus lean paychecks. *Health Economics*, 18: 535–548.
- Harper, B. (2000). Beauty, stature and the labour market: a British Cohort Study. *Oxford Bulletin of Economics and Statistics*, 62: 771–800.
- Heckman, J. (1998). Detecting discrimination. *Journal of Economic Perspectives*, 12: 101–116.
- Henderson, D., S. Kumbhakar and C. Parmeter. (2012). A simple method to visualize results in nonlinear regression models. *Economics Letters*, 17: 578–581.
- Hildebrand, V. and P. Van Kerm. (2010). Body size and wages in Europe: a semi-parametric analysis. CEPS working paper, Luxembourg.
- Hübler, O. (2009). The nonlinear link between height and wages in Germany, 1985–2004. *Economics and Human Biology*, 7: 191–199.
- Ichimura, H. and P. Todd. (2007). Implementing nonparametric and semiparametric Estimators. *Handbook of Econometrics*, Vol. 6B, pp. 5369–5468, Heckman, J. and Leamer, E. (eds.), North-Holland: Amsterdam.
- Jones, F. and J. McMillan. (2001). Scoring occupational categories for social research: a review of current practice, with Australian examples. *Work, Employment and Society*, 15: 539–63.

- Kan, K. and M-J. Lee. (2012). Lose weight for a raise only if overweight: marginal integration for semi-linear panel models. *Journal of Applied Econometrics*, 27: 666–685.
- Kauermann, G., T. Krivobokova and L. Fahrmeir. (2009). Some asymptotic results on generalized penalized spline smoothing. *Journal of the Royal Statistical Society, Series B*, 71: 487–503.
- Kortt, M. and A. Leigh. (2010). Does size matter in Australia? *Economic Record*, 86: 71–83.
- Krivobokova, T. and G. Kauermann. (2007). A note on penalized spline smoothing with correlated errors. *Journal of the American Statistical Association*, 102: 1328–1337.
- Lahey, J. (2008). Age, Women, and Hiring. *Journal of Human Resources*, 43: 30–56.
- Lamina, C., G. Sturm, B. Kollerits and F. Kronenberg. (2012). Visualizing interaction effects: a proposal for presentation and interpretation. *Journal of Clinical Epidemiology*, 65: 855–862.
- Langlois, J., L. Kalakanis, A. Rubenstein, A. Larson, M. Hallam and M. Smoot. (2000). Maxims or myths of beauty? A meta-analytic and theoretical review. *Psychological Bulletin*, 126: 390–423.
- Lee, S. and K. Ryu. (2012). Plastic surgery: investment in human capital or consumption? *Journal of Human Capital*, 6: 224–250.
- Li, Q. and J. Racine. (2007). *Nonparametric econometrics: theory and practice*. New Jersey: Princeton University Press.
- Li, Y and D. Ruppert. (2008). On the asymptotics of penalized splines. *Biometrika*, 95: 415–436.
- Loh, E. (1993). The economic effects of physical appearance. *Social Science Quarterly*, 74: 420–438.
- Maisey, D., E. Vale, P. Cornelissen and M. Tovée. (1999). Characteristics of male attractiveness for women. *Lancet*, 353: 1500.
- Marx, B. and P. Eilers. (1998). Direct generalized additive modelling with penalized likelihood. *Computational Statistics and Data Analysis*, 28: 193–209.
- Mobius, M. and T. Rosenblatt. (2006). Why beauty matters. *American Economic Review*, 96: 222–235.
- Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica*, 46: 69–85.
- Murphy, K. and F. Welch. (1990). Empirical age-earnings profiles. *Journal of Labor Economics*, 8: 202–229.
- Neal, D. and W. Johnson. (1996). The role of premarket factors in black-white wage differences. *Journal of Political Economy*, 104: 869–895.

- Opsomer, J., Y. Wang, and Y. Yang. (2001). Nonparametric regression with correlated errors. *Statistical Science*, 16: 134–153.
- Pagan, J. and A. Davila. (1997). Obesity, occupational attainment and earnings. *Social Science Quarterly*, 78: 756–771.
- Persico, N., A., Postlewaite and D. Silverman. (2004). The effect of adolescent experience on labor market outcomes: the case of height. *Journal of Political Economy*, 112: 1019–53.
- Porter, N. (2003). Sex plus age discrimination: Protecting older women workers. *Denver University Law Review*, 81: 79-111.
- Register, C.A. and Williams, D.R. (1990). Wage effects of obesity among young workers. *Social Science Quarterly*, 71: 130–141.
- Rhode, D. (2010). *The beauty bias: the injustice of appearance in life and law*. Oxford University Press: New York.
- Rooth, D. (2009). Obesity, attractiveness, and differential treatment in hiring: a field experiment. *Journal of Human Resources*, 44: 710–735.
- Ruppert, D. (2002). Selecting the number of knots for penalized splines. *Journal of Computational and Graphical Statistics*, 11: 735–757.
- Ruppert, D., M. Wand and R. Carroll. (2003). *Semiparametric Regression*. Cambridge: Cambridge University Press.
- Sargent, J. and D. Blanchflower. (1994). Obesity and stature in adolescence and earnings in young adulthood: analysis of a British birth cohort. *Archives of Pediatrics and Adolescent Medicine*, 148: 681–687.
- Shimokawa, S. (2008). The labor market impact of body weight in China: a semiparametric analysis. *Applied Economics*, 40: 949–968.
- Singh, D. (1993). Adaptive significance of female physical attractiveness: role of waist-to-hip ratio. *Journal of Personality and Social Psychology*, 65: 293–307.
- Strasak, A., N. Umlauf, R. Pfeiffer and S. Lang. (2011). Comparing penalized splines and fractional polynomials for flexible modelling of the effects of continuous predictor variables. *Computational Statistics and Data Analysis*, 55: 1540–1551.
- Tovée, M., S., Reinhardt, J. Emery and P. Cornelissen. (1998). Optimal BMI and maximum sexual attractiveness. *Lancet*, 352: 548.
- Tovée, M., D. Maisey, J. Emery and P. Cornelissen. (1999). Visual cues to female physical attractiveness. *Proceedings of the Royal Society London, Series B*, 266: 211–218.
- von Hinke Kessler Scholder, S., G. Davey Smith, D.A. Lawlor, C. Propper and F. Windmeijer 2013. Child height, health and human capital: evidence using genetic markers. *European Economic Review*, 57: 1–22.

- Wand, M. (2003). Smoothing and mixed models. *Computational Statistics*, 18: 223–249.
- Wahba, G. (1990). *Spline models for observational data*. Philadelphia: Society for Industrial and Applied Mathematics.
- Wood, S. (2003). Thin plate regression splines. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65: 95–114.
- Wood, S. (2006a). Low-rank scale-invariant tensor product smooths for generalized additive mixed models. *Biometrics*, 62: 1025–1036.
- Wood, S. (2006b). *Generalized additive models: An introduction with R*. Boca Raton, FL: Chapman & Hall.
- Wooden, M., N. Watson, P. Aguis and S. Freidin. (2008). Assessing the quality of the height and weight data in the HILDA survey, HILDA Project Technical Paper Series No. 1/08. Melbourne Institute of Applied Economic and Social Research, University of Melbourne.
- Wooldridge, J. (2010). *Econometric analysis of cross section and panel data*. Second edition. MIT Press: Cambridge.
- Wu, X. and R. Sickles. (2013). Semiparametric estimations under shape constraints. Texas A&M University. Unpublished manuscript.
- Wunder, C., A. Wiencierz, J. Schwarze and H. Küchenhoff. (2013). Well-being over the life span: semiparametric evidence from British and German longitudinal data. *Review of Economics and Statistics*, 95: 154–167.

Table 1: Descriptive Statistics

	Males (<i>n</i> = 13803)		Females (<i>n</i> = 11944)	
	Mean	Std. Dev.	Mean	Std. Dev.
Log weekly earnings	7.19	0.50	6.86	0.48
Weekly earnings	1508.20	857.65	1070.74	556.72
Height (cm)	178.51	7.05	164.60	7.07
Weight (kg)	87.07	15.55	71.83	16.52
BMI	27.30	4.48	26.50	5.84
Underweight	0.004	0.06	0.02	0.14
Healthy weight	0.32	0.47	0.47	0.50
Overweight	0.45	0.50	0.29	0.45
Obese	0.23	0.42	0.22	0.42
Hours of work (all jobs)	44.62	9.49	37.27	9.55
Age	39.41	8.55	40.08	8.74
Experience	9.63	8.78	8.67	8.49
Single	0.16	0.37	0.15	0.36
Couple	0.78	0.42	0.73	0.45
Divorced/Separated/Widowed	0.06	0.24	0.12	0.33
Number of own children	1.01	1.14	1.02	1.11
Number of own children aged 0-4	0.29	0.61	0.14	0.41
University degree	0.30	0.46	0.40	0.49
Diploma	0.10	0.30	0.12	0.32
Certificate	0.31	0.46	0.17	0.37
High school	0.12	0.33	0.14	0.34
High school dropout	0.17	0.37	0.18	0.39
Born in Australia	0.80	0.40	0.80	0.40
Immigrant from English speaking country	0.10	0.30	0.09	0.28
Immigrant from Non-English speaking country	0.10	0.30	0.11	0.32
Indigenous	0.02	0.12	0.02	0.13
Father's occupational scale	45.73	23.19	46.87	23.32
Health condition	0.11	0.31	0.12	0.32
Union	0.30	0.46	0.32	0.47
Employment (ongoing)	0.84	0.37	0.78	0.41
Employment (fixed contract)	0.09	0.29	0.11	0.31
Employment (casual contract)	0.07	0.26	0.11	0.31
Firm size (small)	0.62	0.49	0.62	0.49
Firm size (medium)	0.22	0.42	0.21	0.41
Firm size (large)	0.16	0.37	0.17	0.38
Manager	0.16	0.37	0.10	0.31
Professional	0.24	0.43	0.34	0.48
Trades	0.20	0.40	0.04	0.19
Service	0.06	0.24	0.12	0.33
Administrative	0.09	0.29	0.25	0.43
Sales	0.04	0.20	0.07	0.26
Machinery	0.11	0.32	0.01	0.11
Laborer	0.08	0.28	0.05	0.22
Metropolitan region	0.67	0.47	0.67	0.47
Inner region	0.23	0.42	0.23	0.42
Outer region	0.10	0.31	0.10	0.30
NSW	0.28	0.45	0.28	0.45
VIC	0.25	0.43	0.25	0.44
QLD	0.22	0.41	0.23	0.42
SA	0.08	0.27	0.08	0.27
WA	0.09	0.29	0.08	0.27
TAS	0.03	0.18	0.04	0.18
NT	0.01	0.10	0.01	0.10
ACT	0.03	0.17	0.03	0.16

Notes: Non-missing observations for father's occupational scale are $n=12901$ for males and $n=11231$ for females.

Table 2: OLS Results for Males (Dependent Variable: Log Weekly Earnings)

	OLS				Correlated random effects			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Height (m)	0.337*** (4.30)	0.367*** (4.33)			0.291*** (3.75)	0.296*** (3.62)		
Weight (kg)		-0.000 (-0.88)				-0.000 (-0.21)		
BMI			-0.001 (-0.99)				-0.000 (-0.30)	
Underweight				-0.082 (-1.19)				-0.005 (-0.11)
Overweight				0.033*** (2.98)				0.003 (0.38)
Obese				-0.012 (-0.89)				-0.011 (-1.15)
Other controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	13803	13803	13803	13803	13803	13803	13803	13803
<i>R</i> ²	0.469	0.469	0.467	0.469	0.50	0.50	0.499	0.50


Notes: Robust standard errors are clustered at the person level. ***, **, and * denote statistical significance at the 1, 5, and 10 per cent levels, respectively. See text for control variables used. For regressions based on categorical weight indicators, healthy weight is the excluded category. Wave dummies are included in all models.

Table 3: OLS Results for Females (Dependent Variable: Log Weekly Earnings)

	OLS				Correlated random effects			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Height (m)	0.237*** (3.74)	0.273*** (4.14)			0.202*** (3.10)	0.242*** (3.53)		
Weight (kg)		-0.001* (-1.95)				-0.001** (-2.27)		
BMI			-0.002** (-2.38)				-0.002*** (-2.61)	
Underweight				-0.015 (-0.58)				0.017 (0.77)
Overweight				0.000 (0.04)				0.007 (0.88)
Obese				-0.020* (-1.88)				-0.011 (-1.06)
Other controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	11944	11944	11944	11944	11944	11944	11944	11944
<i>R</i> ²	0.602	0.602	0.601	0.601	0.628	0.628	0.628	0.628

Notes: Robust standard errors are clustered at the person level. ***, **, and * denote statistical significance at the 1, 5, and 10 per cent levels, respectively. See text for control variables used. For regressions based on categorical weight indicators, healthy weight is the excluded category. Wave dummies are included in all models.

Figure 1: BMI, Height and Peak Wages



Notes: BMI values are denoted with numbers accompanying each iso-contour line. The two thick BMI lines (BMI = 23 and BMI = 27) are the approximate BMI values that Gregory and Ruhm (2011) suggest are associated with the peak estimated wages for women and men respectively. The two vertical lines (height = 160 cm and height = 191 cm) are the point estimates of height that Hübler (2009) reports are linked to the highest wages for women and men respectively.

Figure 2: Beauty in Add Health Wave 4 (Males Aged 24-32)




Figure 3: Beauty in Add Health Wave 4 (Females Aged 24-32)




Figure 4: Kernel Weighted Local Polynomial Smooths of Height and BMI





Figure 5: Males – Full Sample ($n = 13803$)

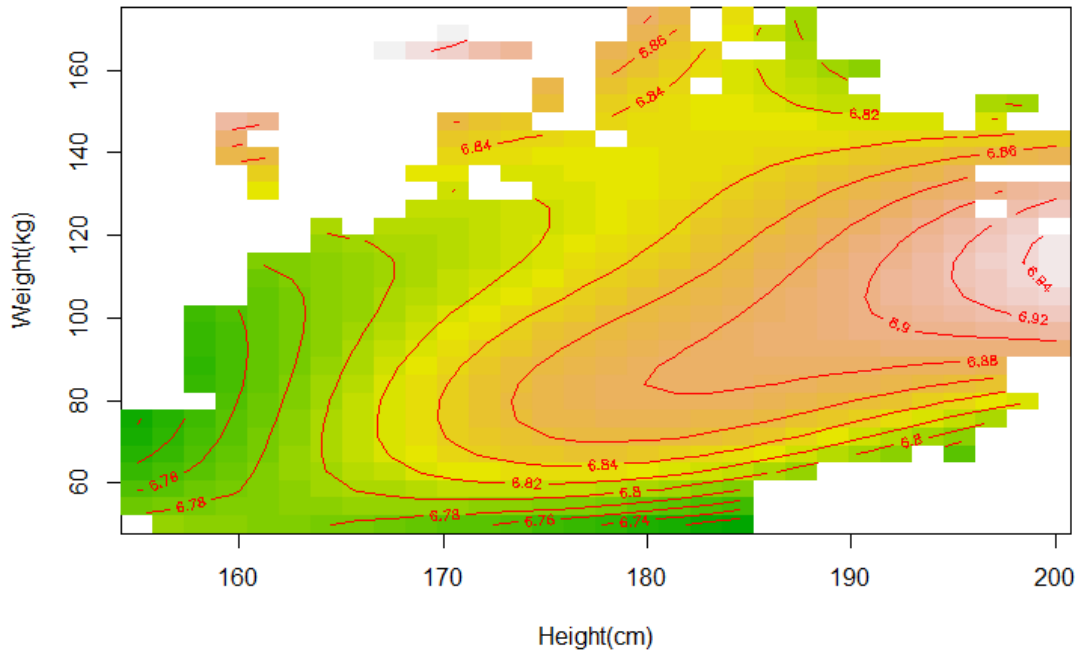


Figure 6: Females – Full Sample ($n = 11944$)




Figure 7: Males Aged 25-39 ($n = 6935$)

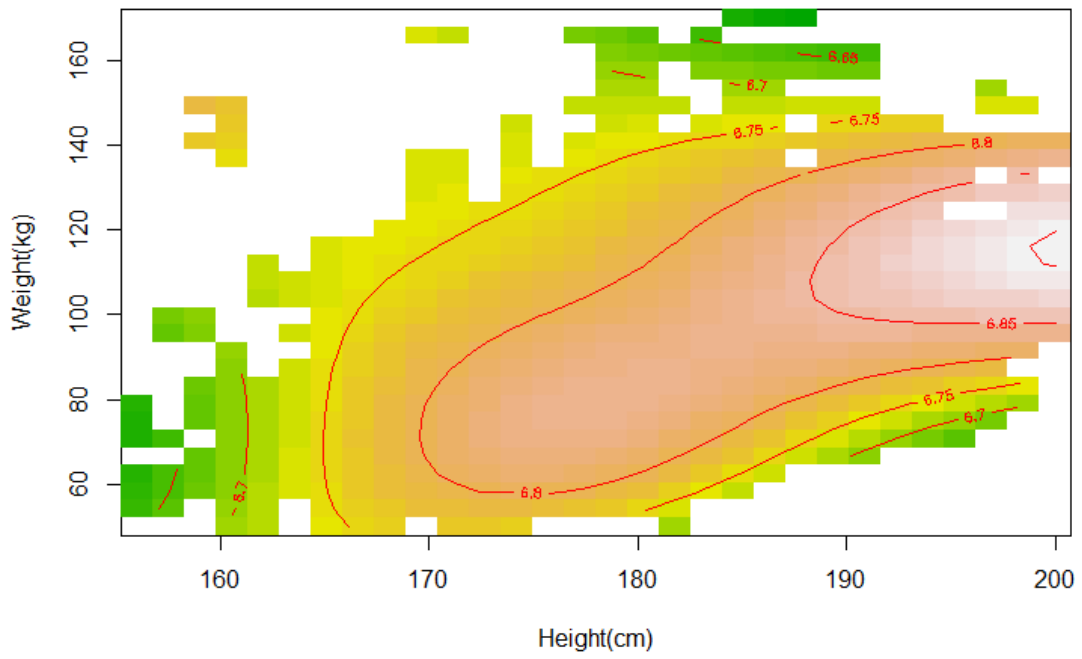


Figure 8: Males Aged 40-54 ($n = 6868$)




Figure 9: Females Aged 25-39 ($n = 5446$)




Figure 10: Females Aged 40-54 ($n = 6498$)

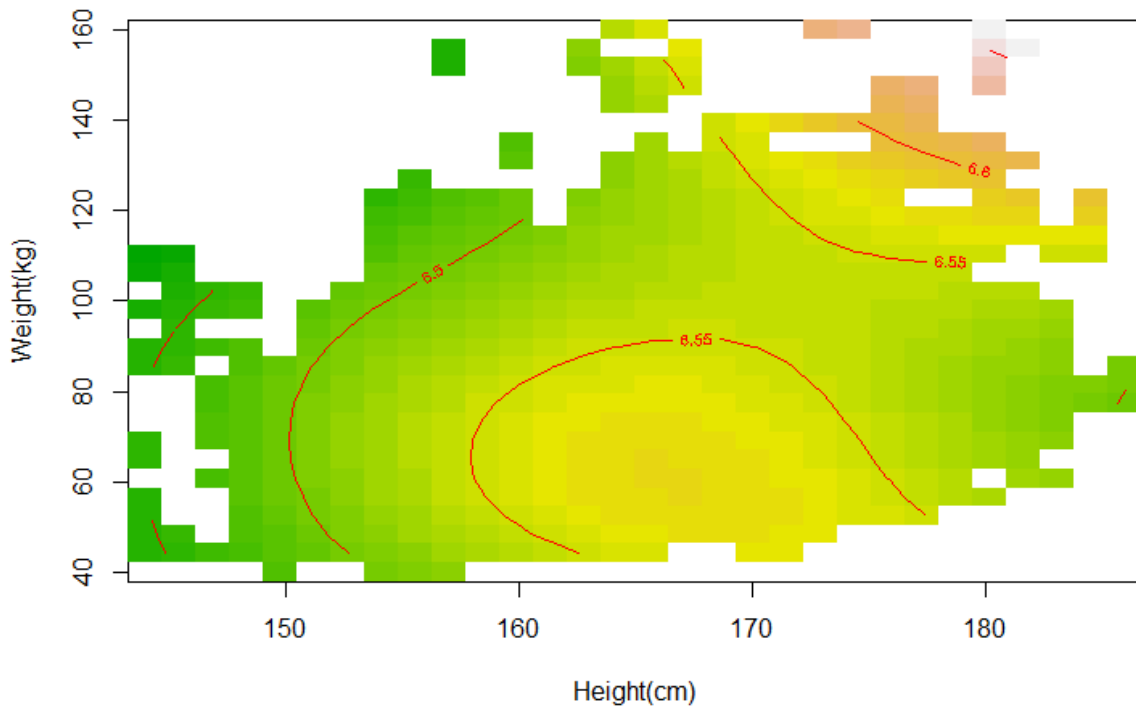


Figure 11: Males Aged 25-54 – Height and Weight in 2006 and Wages in 2012 ($n = 1362$)





Figure 12: Females Aged 25-54 – Height and Weight in 2006 and Wages in 2012 ($n = 1225$)



Appendix

Figure A.1: Area of Thick Support for Height and Weight Combinations in Different Samples

