Employment Effect of Innovation

d'Artis Kancs and Boriss Siliverstovs

1 Introduction

In setting the Research and Innovation (R&I) Strategy 2020–2024, the European Union (EU) has defined six ambitious objectives; including jobs and employment. Research and innovation policy should play a key role in responding to the challenges brought about by the global COVID-19 pandemic. It should help deliver Europe's recovery plan, paving the way out of the current crisis on the path to a fairer future, based on economic growth that places the wellbeing of workers at the centre of the production process (Ivanova et al., 2019).

In the context of these R&I Strategy's objectives, an important policy question arises whether innovation and employment processes can be

d'A. Kancs (✉)
Joint Research Centre, European Commission, Ispra, Italy

LICOS, University of Leuven, Leuven, Belgium
e-mail: d'artis.kancs@ec.europa.eu

B. Siliverstovs
Bank of Latvia, Riga, Latvia

KOF Swiss Economic Institute, ETH Zurich, Zurich, Switzerland

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2022

P. Goulart et al. (eds.), *Global Labour in Distress, Volume I*, Palgrave Readers in Economics,
https://doi.org/10.1007/978-3-030-89258-6_15

complementary and hence their EU targets can be achieved at the same time? Further, policy makers are interested to know: (i) are there R&D intensity levels when innovation and employment are positively related to each other and when innovation may have an adverse impact on the firm employment? (ii) What type of innovators create most jobs and hence provide the highest potential for policy synergies? Answering these questions is the main objective of this study, as they may help to design policies, which can efficiently contribute to achieving both the innovation and employment targets of the R&I Strategy at the same time.

At a first glance, a simultaneous boosting of both employment and innovation may seem an easy and most natural task to achieve as any type of investments (including R&D) increases the labour demand, at least in the short-run. However, the theoretical literature suggests that the relationship between innovation and employment seems to be far more complicated than one can naively assume initially (Smolny, 1998). Also the econometric results reported in the literature on employment effects of innovation are rather contradictory both with respect to their sign and magnitude, suggesting that increasing the innovation intensity can have not only complementary but also substitutionary effects on employment (Young, 1993; Antonucci & Pianta, 2002; Van Reenen, 1997).

In order to accommodate a wide range of possibilities in the innovation-employment relationship ranging from highly negative to strongly positive, in this study we propose an alternative methodological approach that has not been employed in the innovation-employment literature before. In particular, we relax the linearity assumption in the functional relationship between innovation and employment and hope that it will contribute towards sorting out the likely reasons for observing such a large range of estimated employment elasticities with respect to the firm innovation activity. There are several reasons why the innovation-employment relationship may be non-linear. Conceptually, the non-linearities in the functional relationship between innovation and employment may arise, for example, due to the coexistence of many mutually interdependent transmission mechanisms and general equilibrium feedback loops, as the employment effect of innovation depends, among others, on the nature of innovation (product or process innovation); the purpose of innovation (to save labour or capital, neutral, or biased towards skills) and other factors (Pianta, 2004). Empirically, the employment effect of innovation depends on the firm's sector of activity; formal and informal institutions; the time frame of analysis; specifics of the existing production technology;

dimensions of innovation (radical or incremental); consumer preferences; the fierceness of competition in intermediate input and labour markets; the structure of workforce skills; and so on which all contribute to differentiated employment effects at different innovation intensities (Bogliacino & Vivarelli, 2012; Bogliacino et al., 2012; Lachenmaier & Rottmann, 2007).

If the functional relationship between innovation and employment would indeed be non-linear—a fact confirmed in our econometric analysis—then an accurate estimation of the functional relationship would depend crucially on the ability to account for these non-linearities in the innovation-employment nexus, which is highly challenging. Due to complexities related to a suitable counterfactual at the firm level and methodological challenges in the estimation approach, however, there are no studies available in the literature yet that would attempt to account for non-linearities in the R&D and firm employment relationship in a continuous non-linear setting. This study attempts to fill this research gap and estimate the full functional relationship between the firm's innovation and employment in a continuous setup.

To achieve this objective, we rely on a flexible semi-parametric method—the generalised propensity score (GPS) estimator—suggested by Hirano and Imbens (2004). Two main features of the GPS methodology make it particularly attractive for our purpose: (i) estimation can be based on a flexible semi-parametric regression allowing for a non-linear dependence between the variables of interest without imposing any a priori restrictions; and (ii) the elimination of the selection bias arising from a non-random assignment of treatment (R&D expenditure) intensity across firms by conditioning on the observed firm characteristics. In applying the GPS methodology, we attempt to identify the R&D intensity levels under which innovation can be complementary to employment and under which it may have an adverse impact on employment. To the best of our knowledge, the application of a flexible semi-parametric counterfactual methods to the employment-innovation nexus is the first of this sort in literature and hence constitutes our main contribution to literature.

We base our micro-econometric analysis on a large international firm-level panel data set for OECD countries and our proxy for technology is a measurable and continuous variable, while most of previous studies have relied on either indirect proxies of the technological change or dummy variables (such as the occurrence of product and process innovation). In particular, we employ the EU Industrial R&D Investment Scoreboard data set, which comprises data of the R&D investment, as well as other financial and economic variables for the top 2500 innovators worldwide. In addition to firm-level R&D expenditures, we make use also of other economic and financial variables, which allow us to control for important firm-specific characteristics. Moreover, the Scoreboard data also allow to identify the industrial sector (of the parent subsidiary) as well as the geographical region of the R&D investment (according to the location of the firm's headquarter), which allows us to control for fixed sector-specific and location-specific effects.

Our results enhance previous findings by facilitating to connect dots of existing point estimates in literature. Our findings confirm that the relationship between innovation and employment entails important non-linearities. There is notable difference in reaction of employment to the innovation activity of the firm, depending on the actual level of the R&D intensity. It is also worthwhile mentioning that our results also remind that the innovation impact on employment can be negative too—findings that have been reported also in previous studies (Pianta, 2004). For example, in our sample this is the case for companies operating in high-tech sectors, characterising by comparatively high levels of the innovation activity. These results imply that a further increase in R&D expenditures in high-tech sectors can have a non-negligible labour-saving effect. Furthermore, we find that the labour-saving effect of innovation could also be detected for companies operating in low- and medium-low-tech sectors, though this effect is much less pronounced than for highly innovative firms.

The rest of this paper is organised as follows. Next section contains a review of the relevant literature. In Sect. 3 we describe the econometric methodology. The data is described in Sect. 4. The empirical results are presented in Sect. 5. The final section contains conclusions and sets an outline for future research agenda.

2 Previous Literature

The question of whether the technological change creates or destroys jobs has been posed since the beginning of the classical economics of Karl Marx (1867):

> "*Suppose that the making of the new machinery affords employment to a greater number of mechanics, can that be called compensation to the carpet makers, thrown on the streets?*" (Marx (1867): 479).

Ciriaci et al. (2013), Bogliacino et al. (2012), Bogliacino and Vivarelli (2012) and Bogliacino (2014) were among first attempts to decompose the employment effect of innovation according to R&D intensity levels. Using a balanced panel comprising 3300 Spanish firms observed of the period 2002–2009, Ciriaci et al. (2013) investigated the employment effect of innovation both for innovative and non-innovative firms. Ciriaci et al. (2013) found that those firms, which engage more intensively in innovation activities, create more jobs than less innovative firms. In particular, this effect is more pronounced for small and young innovative firms. At the same time they pointed out that for this group of firms, a successful launch of new products in the market as a result of boosting the innovation activity can lead to a higher growth in sales rather than in employment, which is consistent with the labour-saving effects of technological advances, discussed above.

Bogliacino et al. (2012) studied the employment effect of R&D expenditure using the sample of 677 EU firms observed during the period 1990-2008. Employment elasticities were estimated using a dynamic panel model allowing for lagged employment by means of the Least Squares Dummy Variable Corrected (LSDVC) estimator (Bun & Kiviet, 2003; Bruno, 2005). The results were obtained for the sample of all firms as well as for sub-samples comprising service-sector firms, all manufacturing firms and sub-samples comprising manufacturing firms further subdivided into high-tech and non-high-tech firms. The estimated short-run elasticities were 0.023% for the whole sample, 0.056% for service-sector firms and 0.049% for high-tech manufacturing firms. Interestingly, also the corresponding elasticity estimate for non-high-tech manufacturing firms was also positive (0.021%), though not statistically significant. Using the estimated coefficient on the lagged employment variable Bogliacino et al. (2012, Table 1) derived long-run employment elasticities. The long-run elasticities of employment calculated for the whole sample were 0.075%, 0.097% for service-sector firms and approximately of equal magnitude of 0.11% both for all manufacturing firms and high-tech manufacturing firms.

Bogliacino and Vivarelli (2012) conducted study on the employment effect of innovation activity using a sample of 2295 firms from 15 European countries available over the period 1996–2005. All main results of this study were reported for a number of dynamic panel data estimators such as random-effects, fixed-effects as well as two versions of the Generalised Method of Moments [GMM-DIF, Arellano & Bond (1991)] and

[GMM-SYS, Blundell & Bond (1998)], where the last estimator could be identified as the most reliable one (Bogliacino and Vivarelli, 2012, Section IV). These estimators were applied for the whole sample of firms. The short-run elasticity reported by the GMM-SYS estimator was 0.025%, which was very similar to that reported in Bogliacino et al. (2012). However, the long-run elasticity was about 0.31%, which was about four times larger than that reported in Bogliacino et al. (2012) for the whole sample (0.075%). In order to ensure robustness of estimation results, a distinction was made between firms with different levels of the technological sophistication, by allowing for differential employment effects of high-tech, medium-tech and low-tech firms. Employment elasticities were obtained by means of the LSDVC rather than the GMM estimator; as the former estimator outperformed the latter one under given estimation conditions. The main result of Bogliacino and Vivarelli (2012) was that the job creation effect of the R&D expenditure only was evident for the high-tech sector; both for medium- and low-tech sectors the estimated short-run elasticities were not significantly different from zero. For the high-tech sector, short- and long-run elasticities were 0.017% and 0.17%, respectively.

3 Econometric Strategy

In light of the diversity in the channels of adjustment and the reverse causality of interdependencies between innovation and employment, the existing evidence discussed in Sect. 2 suggests that very likely the functional relationship between these two processes is more nuanced than point estimates from previous studies are able to tell us. This implies that an accurate estimation of the functional relationship depends crucially on the ability to account for potential non-linearities in the innovation-employment nexus. In order to allow for a differentiated impact of innovation on employment while accounting for differences among firms at different R&D intensity levels, an appropriate estimation approach is required which does not average across all innovators and employers, but instead allows for a differentiated employment effect at various R&D intensity levels.

To estimate the full functional relationship between innovation and employment, we rely on the generalised propensity score (GPS) approach introduced in Hirano and Imbens (2004).[1] The GPS approach is a further

[1] This approach was already applied to the following pairs of variables: R&D intensity and productivity in Kancs and Siliverstovs (2016), migration and trade in Egger et al. (2012) and

elaboration on the popular binary treatment propensity score estimator of Rosenbaum and Rubin (1983) widely used for impact evaluations of various programmes.[2] In the context of this study, the relevant features of the GPS methodology are as follows. First, it allows for continuous rather than binary treatment levels. Second, it allows to estimate the treatment effect also without a "zero" control group. Third, the GPS procedure eliminates the selection bias arising due to a non-random assignment (choice) of treatment (R&D) intensity across firms by conditioning on observed firm characteristics. Finally, it captures potential non-linearities in the functional relationship between the R&D investment and firm employment, as it relies on a flexible semi-parametric specification.[3] As result, the estimated dose-response functions allow to retrieve the entire interval of average and marginal treatment effects over all possible treatment levels (R&D intensity).

The counterfactual framework of the dose-response analysis naturally involves a dose or treatment variable—R&D intensity—and a response variable—employment—both observed for firm *i*. The difference between usual analysis, typically based on the OLS regression of the response variable on the treatment variable, is that one introduces an additional auxiliary variable, called the generalised propensity score, when modelling the dose-response relationship between the variables of interest. The generalised propensity score is derived from a vector of observed covariates for firm *i*, X_i, and its primary purpose is to remove estimation and inference biases related to non-random dose assignment in the data sample, as discussed above.

Application of the GPS methodology in order to estimate the dose-response analysis typically involves the following three steps (Hirano & Imbens, 2004) . In the first step the GPS variable is constructed using the OLS regression of the treatment variable, r_i, or, as most often in literature, its logarithmic transformation, $\ln r_i$, on a vector of continuous and categorical covariates, X_i, characterising each firm *i* in the data set:

growth effects of the regional policy in the European Union in Becker et al. (2012), *inter alia*.

[2] For an accessible presentation of the logic underlying the propensity-score matching, see Heinrich et al. (2010).

[3] According to Bia et al. (2011), the estimated dose-response function is robust to the choice of a semi-parametric approach, but it is sensitive to a parametric specification.

$$\ln r_i = X_i' \gamma + \varepsilon_i, \qquad \varepsilon_i \sim N(0, \sigma^2). \tag{1}$$

Observe that a usual assumption that is made is that the distribution of the error terms is normal with variance σ^2. If this assumption is supported by the data then the GPS variable is defined as the normal probability density function estimated for the regression residuals:

$$\hat{s}_i = \frac{1}{\sqrt{2\pi\hat{\sigma}^2}} \exp\left[-\frac{1}{2\hat{\sigma}^2} \hat{\varepsilon}_i^{\,2} \right]. \tag{2}$$

Hirano and Imbens (2004, Section 7.4) mention that other more flexible distributions can be used in case if the normality assumption is not supported by the data, for example, this could be a mixture of several normal distributions or a normal heteroscedastic distribution, in which variance is a function of covariates. Alternatively, the departures from normality can be accommodated by non-parametric approach that relies on the kernel probability density estimation. In the empirical part of this paper, we resort to the latter option, since we find that the normality assumption is violated in our sample. To the best of our knowledge, this is the first attempt to rely on non-parametric methods for estimation of the generalised propensity score in the literature and it serves as an additional methodological contribution to the relevant literature.

The propensity score in Eq. (2) fulfils its purpose of measuring the degree of similarity across heterogeneous firms when the so-called *balancing* property is satisfied, that is, for those firms with assigned equal propensity scores (conditional on firm-specific covariates) the associated treatment level is independent of firm characteristics. In this step, we follow the test procedures of Hirano and Imbens (2004) in order to verify whether the balancing property is not violated in our data sample.

In the second step, the expected value of response variable, $\ln \omega_i$, is modelled as a flexible semi-parametric function of the treatment variable and the estimated generalised propensity score, $\ln r_i$ and s_i, respectively:

$$\begin{aligned} E[\ln \omega_i \mid \ln r_i, s_i] = Incpt &+ \alpha_{11} * \ln r_i + \alpha_{12} * \left[\ln r_i\right]^2 + \alpha_{13} * \left[\ln r_i\right]^3 \\ &+ \alpha_{21} * s_i + \alpha_{22} * \left[s_i\right]^2 + \alpha_{23} * \left[s_i\right]^3 \\ &+ \alpha_3 * (\ln r_i * s_i), \end{aligned} \tag{3}$$

where the latter variable is substituted with its estimates, $\hat{s}_i$, from the first step. The flexibility of the functional form can be controlled for by varying the power of variables $\ln r_i$ and s_i and their cross-products.

The average expected response of the response variable, ω, for a given treatment dose, ρ, is estimated in the third step:

$$E[\ln\hat{\omega}(\ln\rho)] = \frac{1}{N}\sum_{i=1}^{N}\Big[Incpt + \hat{\alpha}_{11} * \ln\rho + \hat{\alpha}_{12} * [\ln\rho]^2 + \hat{\alpha}_{13} * [\ln\rho]^3 \\ + \hat{\alpha}_{31} * \hat{s}(\ln\rho, X_i) + \hat{\alpha}_{32} * [\hat{s}(\ln\rho, X_i)]^2 + \hat{\alpha}_{33} * [\hat{s}(\ln\rho, X_i)]^3 \\ + \hat{\alpha}_3 * (\ln\rho * \hat{s}(\ln\rho, X_i))\Big], \tag{4}$$

where the coefficient estimates from Eq. (3) are used. The whole dose-response function is obtained by computing Eq. (4) for each treatment level by using a grid of values in the corresponding range of the treatment variable.

In the final step, we derive the treatment effect function as a first derivative of $E[\mathrm{I}\hat{\omega}(\mathrm{In}\rho)]$ with respect to argument $\ln\rho$. By definition the treatment effect function computed in this way measures estimated employment elasticity with respect to R&D, allowing us to directly compare our results with those reported in the existing literature. Following Hirano and Imbens (2004), confidence intervals around the estimated dose-response and treatment effect functions are obtained by means of a bootstrap procedure.

4 Data Sources, Sample and Variable Construction

Data Sources

The principal data source is the EU Industrial R&D Investment Scoreboard maintained by the European Commission. The R&D Scoreboard is an annual data set that comprises firm-level data on the R&D investment, as well as other financial and economic variables (e.g. net sales, operating profits, employees) for the top 2500 R&D performers worldwide. In addition to economic and financial variables, the R&D Scoreboard also identifies the main industrial sector (of the parent company) as well as the geographical region of R&D investment (according to the location of company's head-quarter).

An important limitation of the R&D Scoreboard data concerns the issue of non-random sample selection, putting under question the general validity of our results. Given the underlying sampling and selection rules of the R&D Scoreboard data set—ranking and selecting companies according to the total amount of their R&D expenditures—the R&D Scoreboard is not a random sample. Hence the R&D Scoreboard data set may be criticised that it has a sample bias affecting the results, as it only represents top R&D investors. However, given our interest in the employment effect of innovation, this issue is of lower order of magnitude, because we are covering almost the entire population of the world-wide R&D investment (Moncada-Paterno-Castello et al., 2010). As described below, out of the 2500 firms listed in the R&D Scoreboard data only for 1659 companies there were complete data records, prompting us to analyse the available data.[4] Still, these 1659 Scoreboard's companies selected for this study represent around 80% of the world-wide business R&D expenditure. While small R&D investors and non-R&D-performers are excluded from our sample, the aim of this study is to focus on the impact of the R&D-driven innovation on employment, but not to examine determinants of the labour demand in the entire economy. Finally, the particular estimation approach that we adopt in this study allows us to estimate counterfactual treatment effects also without a zero control group.

Sample Construction

In this study, we use R&D Scoreboard data for the last four available years: 2014-2017. Our choice of this sample is motivated by the fact that it is a reasonably long period apart from the Great Financial Crisis (GFC) that undoubtedly had pronounced effects on the firms' investment activity. Including observations from years during the GFC and shortly after its outbreak had a distortive impact on the long-run relationship between innovation and employment prevailing in the business-as-usual environment that we aim to capture in our study.

Since the Scoreboard involves individual firm-level data covering many countries, industries and technological levels of sophistication, it is rather unsurprising that due to all this incumbent data heterogeneity the

[4] Companies which do not disclose figures for R&D investment or which disclose only figures which are not material enough were also omitted from our analysis.

annual data for top 2500 Scoreboard companies forms an unbalanced panel. There are firms that were not present among the top 2500 R&D performers either in the beginning or in the end of the sample period or even at the both ends of the sample period and hence have missing observations. There are also firms that were present in the top 2500 Scoreboard sample at the beginning and at the end but have missing data points for some years within our sample period. All this implies a loss of observations, if our identification strategy aimed at exploiting both inter-temporal and cross-sectional dimensions of a balanced panel. Another option would be to focus solely on the cross-sectional dimension for a particular year, but this again involves loss of information as well as a certain arbitrariness in the choice of the particular year. Hence, in order to retain as many observations as possible, we construct our sample from firms for which there are at least two consecutive years of observations for all variables of our interest. For these firms, we compute averages of their characteristics using the available observations. This helps us smoothing year-on-year fluctuations in our data and avoid a potential source of outlier bias.

Finally, we did a sanity check for the resulting sub-sample of firms and filtered out firms that have extreme values of the R&D intensity which, as discussed in Sect. 4, is defined as the ratio of the R&D investment to net sales. In particular, we removed firms for which the estimated R&D intensity exceeds unity. For this sub-group of firms the median R&D intensity is 6, whereas the maximum is 1210. It turns out that *all* these firms are characterised by a rather small actual employment (the median employment is 113 persons) and a negative operating profit. The former fact indicates that the share and hence the impact of these firms on the total employment are rather small. Moreover, the latter fact indicates that such business model/innovation pattern is not sustainable in the long run. Therefore, in order to make our sample more homogeneous we treat these firms as outliers that need to be removed from the empirical analysis. As a result of data cleaning, we are left with 1659 observations that form the basis for our empirical analysis.

Data Set

The dependent (response) variable is a firm-specific employment measured by the number of employees (EMPL). For each firm in our sample we use the average number of employees for the available years. These companies included in our sample data employed around 44.1 mln. workers with largest shares of about 10.6 and 14.5 mln. workers pertaining to companies registered in the US and the EU. The R&D investment totalled 2028 milliard Euro with about 42% and 28% of the total sum is attributable to the companies from the US and the EU, with the Japan and China accounting for about 17% and 6%, respectively.[5]

The remaining firm characteristics (Net sales (NSALES), Operating profit (OP), Capital expenditure (CAPEX)) contained in the Scoreboard were complemented with Market capitalisation (MCAP) sourced from both the Financial Times London Share Service and Reuters. In order to create a relative measure of R&D expenditure that takes into account firm commercial size, we create the treatment variable (R&D intensity) as the ratio of the nominal R&D expenditure to net sales.

There are several categorical dummy variables indicating level of technological sophistication (low-tech, medium-low tech, medium-high tech and high-tech) that are further sub-divided into industrial sectors according to the ICB classification as well as dummy variables indicating countries. Further details on the definitions of the explanatory variables are provided in the online appendix.

The set of covariates used in our analysis is selected based on previous studies (e.g. see Hall et al., 2008), subject to their availability in our data set. In order to provide an impression on the magnitude of the main firm characteristics and their relationship to the variables of our main interest we report median values of these characteristics evaluated at each level of technological sophistication, see Table 1.

The first observation is that the number of firms belonging either to high- or medium-high-tech sectors (1367) is much larger than the number of firms belonging either to low- or low-tech sectors (292). Such an over-representation of the high-tech firms in the sample naturally reflects the original intention of collecting and maintaining the database on the world top R&D performers. In terms of employment, a median firm-specific employment is inversely proportional to the level of technological sophistication: in the high-tech sector the median employment is 4200 whereas in the low-tech sector it comprises 20,960 employees. In nominal

[5] Note, however, that data reported by the Scoreboard companies do not inform about the actual geographic distribution of the number of employees. A detailed geographic analysis should take into account the location of subsidiaries of the parent Scoreboard companies as well as the location of other production activities involved in the value-chains.

Table 1 Median values for variables by tech sector

Tech sector	*Obs.*	*EMPL (number)*	*R&D (€mln.)*	*R&D sectoral share*[a]	*R&D intensity*	*OP (€mln.)*	*NSALES (€mln.)*	*CAPEX (€mln.)*	*MCAP (€mln.)*
High	649	4200	79.8	0.532	0.115	87.6	907	40.2	2622
Medium-high	718	10,898	63.8	0.363	0.034	184.2	2368	95.6	2724
Medium-low	105	14,073	76.1	0.042	0.016	348.3	4636	206.6	5562
Low	187	20,960	74.0	0.063	0.011	567.4	9589	435.2	5929

[a]Indicates share of each tech sector in total volume of R&D expenditure

terms, the median level of R&D expenditure is about the same across the different tech sectors with a typical value of about 60–80 mln. Euro. However, the sector-specific share of R&D expenditure is not equally distributed as indicated in the column "R&D sectoral share." The lion's share of the total R&D expenditure (about 90%) is accounted for the firms in high- or medium-high-tech sectors.

As far as the treatment variable (R&D intensity) concerns, the median level is highest for the firms in the high-tech sector and it continuously decreases with the level of the technological sophistication. A median firm in the high-tech sector spends about 11.5% of its net sales volume on R&D, whereas the corresponding share for a median firm in the low-tech sector is about 1%.

It is also interesting to observe that the median values of the financial variables like operating profit, net sales, capital expenditure and market capitalisation are highest for the low-tech firms and the lowest for the high-tech firms.

5 Results

This section is sub-divided into two parts. In the first part, we report estimation results from a naive OLS regression of employment on the R&D intensity. Despite the associated econometric issues, this naive model can serve as a useful benchmark against which we can compare the results of more sophisticated methodology based on the generalised propensity score approach applied to the estimation of the functional relationship between the variables of interest, reported in the second part of this section.

OLS Estimation

The scatterplot of employment against R&D intensity is shown in Fig. 1 along with the fitted regression line. The OLS coefficient estimates are shown in the figure as well. The OLS estimate of the employment elasticity with respect to the R&D intensity is reported -0.739 indicating that a 1% increase in the R&D intensity is associated with 0.74% decrease in the number of employees. With the estimated standard error of the slope coefficient 0.026 this elasticity estimate is statistically significantly different from zero and the regression is characterised by a rather goodness of fit with the associated $R^2 = 0.327$.

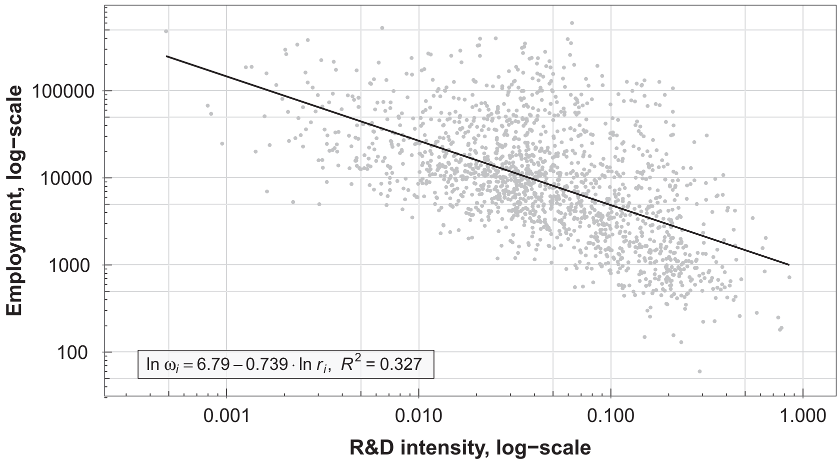

Fig. 1 OLS regression: all firms

GPS Estimation

As explained in Sect. 3 above, the application of the GPS methodology in order to estimate the dose-response function involves three steps. The results of the first step GPS estimation procedure (see Eq. (1)) are reported in Table 2. They suggest that the variation in the R&D intensity is best captured by variables such as the total capital expenditure and its square, market capitalisation and its square, as well as operating profits. Also the included industry- and region-specific dummy variables contribute substantially to the explanatory power of the first step of the GPS regression.[6] Indeed, the goodness-of-fit of this regression is quite high, yielding a R^2 of 68.2%, which is necessary in order to create a mighty propensity score able to remove biases when estimating the dose-response function between the variables of interest.

The assumption of normally distributed OLS residuals in Eq. (1) is verified by means of the Shapiro-Wilk normality test, yielding the p-value of 1.746×10^{-15}. Hence our data do not support the normality assumption. Therefore it is instructive to take a closer look at the histogram of the regression residuals, shown in Fig. 2. The fitted normal probability density function is shown as the dashed line. As seen, the residuals are characterised by too large excess kurtosis and appear to be left-skewed to be compatible with

[6] These are not shown in the regression output table in order to save the space.

Table 2 Dose regression

	Dependent variable: ln *R&D intensity*
$\ln \text{CAPEX}$	-0.326*** (0.052)
$[\ln \text{CAPEX}]^2$	0.185*** (0.057)
$\ln \text{MCAP}$	-0.292** (0.120)
$[\ln \text{MCAP}]^2$	0.315*** (0.078)
$\ln \text{OP}^+$	-0.106* (0.064)
$[\ln \text{OP}^+]^2$	-0.100 (0.072)
$\ln \text{OP}^-$	0.030 (0.084)
$[\ln \text{OP}^-]^2$	0.006 (0.013)
Constant	-1.070* (0.600)
Observations	1659
R^2	0.682

Notes: The sign-preserving log transformation of the operating profit variable was carried out as follows: for positive values $\ln \text{OP}^+ = \text{if } OP > 0 : \ln OP$ and zero otherwise; for negative values $\ln \text{OP}^- = \text{if } OP < 0 : -\ln(\text{-OP})$ and zero otherwise. Sectoral and country dummies (not shown) were included in the regression

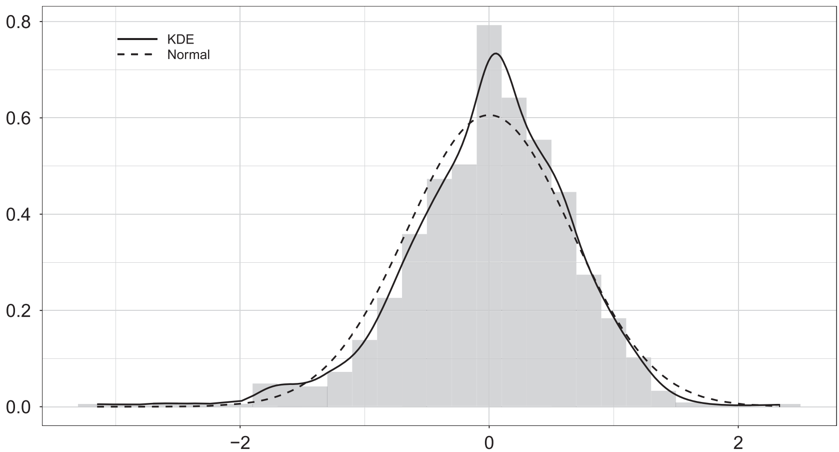

Fig. 2 Equation (1): Residuals histogram; parametric (Normal) and non-parametric kernel density estimation (KDE)

the normal distribution. Hence, instead of relying on the unfulfilled normality assumption, we estimate the GPS by means of non-parametric approach using a kernel density estimation (KDE) of the probability density function, since we have a rather large data set of 1659 observations.

The estimated non-parametric GPS is shown as the solid line in the figure. Due to its inherent flexibility, the kernel-estimated GPS matches the empirical distribution of residuals much better than the one based on the normal distribution. The GPS range is quite large [0.00289, 0.734], signifying substantial differences in the estimated propensity of the treatment level assignment across firms in our sample.

According to the estimation procedure outlined in Hirano and Imbens (2004), the next step is verification of the so-called balancing property of the GPS, preceded by imposing the common- support restriction on the data in question. The latter procedure aims to construct a more homogenised sample by filtering out aberrant observations for which propensity-score-based matching turns out problematic. The former procedure aims at testing whether conditional on observed values of the GPS variable there are no systematic differences in firms' characteristics irrespective of the assigned treatment intensity. As discussed in the online Appendix, the imposition of the common-support restriction reduced the number of firms available for the further analysis from 1659 to 1296. At the same time, the balancing property of the constructed GPS in Eq. (1) is supported by the data, see the Appendix for further details.

Next, we proceed to the estimation of the dose-response relationship between the firm innovation and employment variables. The estimation results for the second-step regression corresponding to Eq. (3) are reported in Table 3. Second step regression results clearly show that the employment response to the firm innovation (proxied by R&D expenditures) is highly non-linear, as all included polynomial terms of the latter variable report highly significant coefficients. It is also worthwhile noticing that the GPS variable enters as a statistically significant covariate both

Table 3 Conditional regression

	Dependent variable: ln *EMPL*
ln R & D intensity	-3.150*** (0.616)
$[\ln \text{R \& D intensity}]^2$	-0.869*** (0.210)
$[\ln \text{R \& D intensity}]^3$	-0.083*** (0.022)
GPS	-3.040*** (0.629)
GPS $*$ln R&D intensity	-0.885*** (0.205)
Constant	5.510*** (0.570)
Observations	1296
R^2	0.226

in levels and via the interaction term with the (log) of our treatment variable, confirming its relevance in eliminating the sample selection bias.[7] The resulting R^2 is 22.6%, which is of a comparable magnitude reported in other studies (Egger et al., 2012).

In order to facilitate the interpretation of the estimation results, we have plotted the estimated dose-response and marginal treatment effect functions in the upper and middle panels of Fig. 3, respectively. The bands around the estimated functions are 95% bootstrap confidence intervals. Observe that in order to facilitate the description of the results in the lower panel of the figure we have plotted the cumulative share of employment in the firms in our sample as a function of the R&D intensity. The curve in the lower panel reveals that 90% of employment in our data sample is accounted by firms with the R&D intensity in the interval between 0.6% and 15%. There are 1088 out of 1296 firms, or about 84%, of the total sample in this interval. There are 21 and 187 firms in the left and right 5% tails of the cumulative employment distribution sorted by the R&D intensity.

The shape of the estimated dose-response function is generally downward sloping, which is broadly consistent with the naive OLS estimation results reported in Sect. 5. However, recall that according to the OLS results the estimated employment elasticity is uniformly negative at all R&D intensity levels. In contrast, the estimated dose-response function using the GPS suggests that the magnitude of the response of employment to changes in the R&D intensity varies with the level of the firms' innovation intensity. This non-linearity in the employment response is well illustrated by the marginal treatment effect function, which can be interpreted as employment elasticity with respect to R&D intensity, that is shown in the middle panel of Fig. 3.

The estimated elasticity of interest has a hump-shaped form. Hence it is convenient to summarise our findings by distinguishing between different treatment intensity levels taking into the consideration such hallmarks as the top and bottom 5% cumulative employment thresholds. For relatively low treatment intensity levels (below 0.6%) the employment elasticity increases in the absolute value from -0.5% up to about -1.5% as the treatment intensity falls. However, given a rather small number of observations

[7] Higher order power transformations of the GPS variable turned out to be insignificant and therefore were omitted from the model specification for the sake of parsimony.

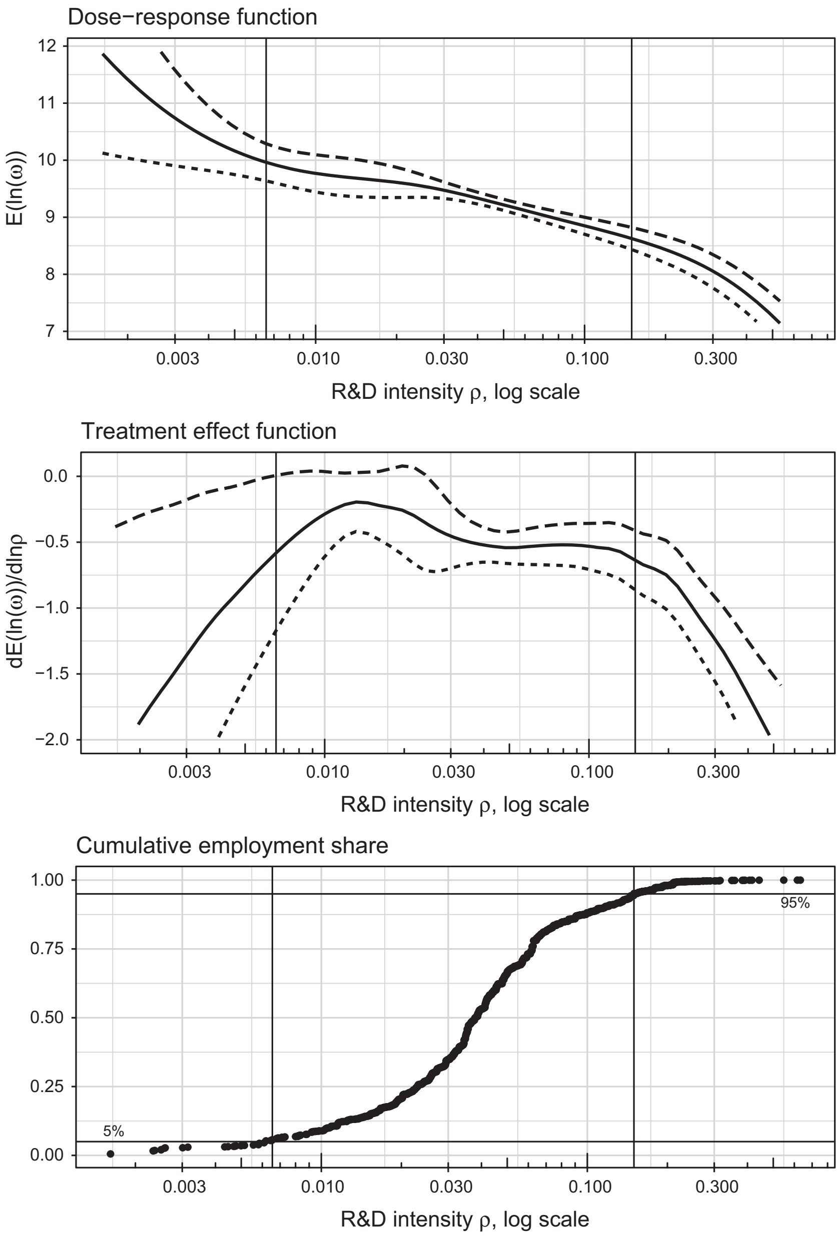

Fig. 3 Dose-response, treatment effect functions and cumulative employment share

in this part of the distribution these estimate values have to be taken with caution.

For the firms within the central 90% interval of the treatment intensity one can make the following two observations. First, for the firms with R&D intensity in the interval between 0.6% and 3% the estimated elasticity is not significantly different from zero, as the bootstrapped 95% confidence interval includes zero line. This suggests a labour-neutral effect of innovation for the firms with medium-low and medium levels of innovation intensity. Second, for the firms with the medium-high levels of the R&D intensity pertaining to the interval between 3% and 15% the estimated elasticity is negative and significantly different from zero. For these firms it is estimated around -0.5% with the associated 95% confidence interval about (-0.3%, -0.7%), suggesting labour-saving effect of innovation. Notwithstanding that this value is substantially lower than that reported by the OLS estimation (-0.74%) earlier in the text.

Turning to the firms with the highest R&D intensity (>15%), this labour-saving effect turns out to be even more pronounced. In this interval, the estimated employment elasticity gradually increases (in the absolute magnitude) from -0.5% to -2.0%, suggesting that the innovation leaders tend to react more and more disproportionately stronger to changes in the R&D intensity in reducing their labour force than innovation followers and moderate innovators.

All in all, our estimation results when compared to those from the naive OLS regression suggest that the employment effect of innovation varies with the level of technological sophistication and warrant against application of estimation techniques that does not accommodate such level dependence. For the firms with rather low to medium ratios of R&D expenditure to net sales this effect tends to be overestimated by the OLS regression whereas understated for the firms on the other side of the spectrum characterised by high values of R&D intensity.

It is instructive to compare our results with traditional point estimates available in the previous literature, despite the fact that studies summarised in Sect. 2 focus on the employment elasticity with respect to a nominal measure of the R&D expenditure, whereas we focus on the employment elasticity with respect to a relative measure of the R&D expenditure. Our results, emphasising the complexity of the non-linear relationship between employment and innovation, are complementing those of Bogliacino (2014), who equally finds that R&D investment expenditures have a non-linear effect on the firm employment, depending on the R&D intensity.

However, compared to the most of the published literature, our results reveal no support for a job-creating aspect of innovation at least when the world top R&D performers are scrutinised. For this particular sub-sample of firms we find that the effect of innovation is at best labour-neutral at the relatively low values of the R&D intensity. For higher innovation intensity levels, the labour-saving effect of innovation becomes increasingly pronounced, as knowledge-intensive firms are looking for high-skilled labour force which is typically in much shorter supply and correspondingly more expensive than their low-skilled fellows.

6 Conclusions, Policy Recommendations and Limitations

The objective of the study is to expose the entire innovation-employment relationship for different R&D intensity levels in a continuous framework. We use a large international firm-level panel data set for OECD countries and employing a flexible semi-parametric method—the generalised propensity score—allows us to estimate the full functional relationship between the R&D-driven innovation and firm employment as well as address important econometric issues, which is not possible in the standard estimation approach used in the previous literature. This is our main contribution to the academic literature and policy debate; to the best of our knowledge no comparable studies analysing the employment effect of innovation in a continuous setting are available in the literature.

In order to answer these questions, we have based our empirical microeconometric analysis on a large international firm-level panel dataset for OECD countries, and our proxy for technology has been a measurable and continuous variable, while the majority of previous studies have relied on either indirect proxies of the technological change or dummy variables (such as the occurrence of the product and process innovation). In particular, we have employed the EU Industrial R&D Investment Scoreboard data set for 2500 R&D performers worldwide. In addition to firm-level innovation expenditures, we have used also economic and financial variables, which allowed us to control for important firm-specific effects, along with sectoral and regional dummies.

Our results suggest that a care should be taken when analysing employment-innovation nexus. Depending on the level of R&D intensity, we find that the innovation impact on employment can be negative too—findings that have been reported also in previous studies. This labour-saving aspect of innovation is more pronounced for firms with medium-high levels of R&D intensity and it tends to increase with the levels of R&D intensity. In terms of policy recommendation, our results imply that these companies should not be immediately targeted by policies aiming to achieve both innovation and employment targets of the R&I Strategy in the same time.

Turning to limitations of our study, an important caveat of our empirical analysis concerns the nature of the Scoreboard sample. First, while other data sets, such as the OECD BERD data, can be considered as fully representative of OECD economies, in the EU Industrial R&D Investment Scoreboard data used in this study only R&D "champions" are considered. This is a clear limitation of our data, the results of which cannot be straightforwardly extrapolated to, for example, SMEs.

A further limitation of the data used in our study is that R&D Scoreboard data do not allow us to identify the effect of product and process innovations separately. However, as discussed in the introduction, the employment effect of innovation can be very different depending on the nature of innovation. In order to separately identify the employment effect of the product and process innovation, other sources of data, such as the Community Innovation Survey (CIS), need to be used, which is a promising area for the future research.

Lastly, in our study we focus on the snapshot of the economy at one period of time without taking higher order effects of firms innovation activity. In the longer run, investing in the innovation activity encourages knowledge-based economy, drives demand for high-skilled, educated workers and eventually brings a country on the higher growth path. However, a comprehensive assessment of these effects is only possible within general-equilibrium models that capture vertical and horizontal linkages between firms, which is not possible to account for in microeconometric studies, such as the one presented in this paper (Kancs & Ciaian, 2011; Brandsma & Kancs, 2016). Hence aligning our results with macro results is indeed important for enhancing our understanding of the employment effect of innovation and it sets a promising avenue for the future research.

Acknowledgments The authors acknowledge helpful comments from Francesco Bogliacino and Daria Ciriaci, Mihails Hazans, Enkelejda Havari, Thomas Hemmelgarn, Giuseppe Piroli, the editor of the journal, two anonymous referees, participants of the conferences on Counterfactual Methods for Policy Impact Evaluation in Rome and Milan, the Scottish Economic Society Annual Conference in Perth and the 57th Meeting of the Italian Economic Association in Milan as well as participants of the research seminars at the European Commission, Bank of Latvia, University of Leuven and Stockholm School of Economics in Riga. We are grateful to Pietro Moncada-Paterno-Castello and the IRI team for granting access to the EU Industrial R&D Investment Scoreboard data. This article follows the methodological framework of Kancs and Siliverstovs (2019). The authors are solely responsible for the content of the paper. The views expressed are purely those of the authors and may not under any circumstances be regarded as stating an official position of the European Commission or the Bank of Latvia. Computations were performed in ***R*** (R Development Core Team, 2008).

Appendix

Explanatory Variables

The following groups of the explanatory variables were used in our analysis in the main text:

- *Net sales*, NSALES: In line with the accounting definition of sales, sales taxes and shares of sales of joint ventures & associates are excluded. For banks, sales are defined as the "Total (operating) income" plus any insurance income. For insurance companies, sales are defined as "Gross premiums written" plus any banking income.
- *Operating profit*, OP: Profit (or loss) before taxation, plus net interest cost (or minus net interest income) and government grants, less gains (or plus losses) arising from the sale/disposal of businesses or fixed assets. Due to the fact that companies report both positive and negative operating profit, we cannot take a logarithmic transformation of this variable. In order to do so, we created the following two variables $\ln \mathrm{OP}^{+}$ and $\ln \mathrm{OP}^{-}$. The former variable is equal to the log of actual values whenever a firm reports positive profit and zero otherwise. The latter variable is equal to the log of absolute actual values multiplied by minus one whenever a firm reports negative profit and zero otherwise.
- *Capital expenditure*, CAPEX: The expenditure used by a company to acquire or upgrade physical assets such as equipment, property, industrial buildings. In company accounts capital expenditure is added to the asset account (i.e. capitalised), thus increasing the

amount of assets. It is disclosed in accounts as additions to tangible fixed assets.

- *Market capitalisation*, MCAP: The share price multiplied by the number of shares issued at a given date. Market capitalisation data have been extracted from both the Financial Times London Share Service and Reuters. These reflect the market capitalisation of each company. The gross market capitalisation amount is used to take into account those companies for which not all the equity is available on the market.
- *Country dummies*: There are 36 distinct countries included in the estimation sample.
- *Industry sector dummies*: The industry sectors are based on the ICB classification. The level of disaggregation is generally the three-digit level of the ICB classification, which is then converted to NACE Rev.2.
- *Sectoral dummies*: In order to account for the sectoral heterogeneity with respect to the R&D intensity, we regroup all firms into four sub-samples according to the level of their technological sophistication. Following the OECD classification, all firms in our sample are regrouped into four 3-digit Industry Classification Benchmark (ICB) groups: high-, medium-high-, medium-low- and low-tech companies:

  – *High-tech*: Technology hardware & equipment, Software & computer services, Pharmaceuticals & biotechnology, Health care equipment & services, and Leisure goods;
  – *Medium-high-tech*: Industrial engineering, Electronic & electrical equipment, General industrials, Automobiles & parts, Personal goods, Other financials, Chemicals, Aerospace & defence, Travel & leisure, Support services, and Household goods & home construction;
  – *Medium-low-tech*: Food producers, Fixed line telecommunications, Beverages, General retailers, Alternative energy, Media, Oil equipment, services & distribution, and Tobacco;
  – *Low-tech*: Gas, water & multi-utilities, Oil & gas producers, Nonlife insurance, Industrial metals & mining, Construction & materials, Food & drug retailers, Banks, Electricity, Industrial transportation, Mobile telecommunications, Forestry & paper, Mining, Life insurance.

Verification of GPS Balancing Property

The balancing property of the constructed GPS variable is verified following the procedure suggested by Hirano and Imbens (2004). Each covariate is subdivided into three groups of 553 observations according to the percentiles of the distribution of the treatment intensity variable. The initial testing of the balancing property amounts to testing whether the average value of a particular variable in every group is equal to the average value in the remaining groups. The results of these tests are reported in Table 4. Only for a handful of covariates we cannot reject the tested null hypothesis at usual significance levels, indicating that there is a strong heterogeneity among covariates belonging to these three groups pertinent to different values of the treatment intensity. A well-specified GPS should be able to successfully account for these differences.

Before verifying the balancing properties of the GPS, we impose the so-called common-support restriction. The purpose of this restriction is to filter out observations that are rather dissimilar in their characteristics when used for the GPS computation in the first step, see Eqs. (1) and (2). As argued by Becker et al. (2012) , it is advisable to impose the common-support condition in order to improve the balancing properties of the GPS and hence achieve more reliable estimation results.

For each treatment group, defined above, $k = 1, 2, 3$, we evaluate GPS values for each observation i at the respective median treatment value, GPS_i^k . We determine the common support region by comparing values of GPS_i^j for each $j = k$ with those computed for other groups $j \neq k$ at the median treatment value of the selected group j. Those observations for which $GPS_i^{j \neq k}$ that fall outside of the range of GPS_i^j are labelled as those

Table 4 Covariate balance, t-statistics (initial data)

	Dose group 1	*Dose group 2*	*Dose group 3*
$\ln \text{CAPEX}$	-18.35	-1.86	19.32
$[\ln \text{CAPEX}]^2$	-14.91	-0.28	17.44
$\ln \text{MCAP}$	-8.93	2.01	6.19
$[\ln \text{MCAP}]^2$	-8.11	2.28	5.44
$\ln \text{OP}^+$	-14.36	-3.91	15.80
$[\ln \text{OP}^+]^2$	-13.13	-1.22	14.03
$\ln \text{OP}^-$	-5.08	-6.13	8.76
$[\ln \text{OP}^-]^2$	-3.17	-4.62	6.69

Table 5 Number of observations in common support

Common support	*Dose group 1*	*Dose group 2*	*Dose group 3*
1296	1651	1430	1411

Table 6 Number of observations by dose group and block

	1:(Total)	*1*	*~1*	*2:(Total)*	*2*	*~2*	*3:(Total)*	*3*	*~3*
m(GPS) = 1	738	68	670	492	105	387	776	87	689
m(GPS) = 2	200	68	132	230	104	126	212	87	125
m(GPS) = 3	146	68	78	207	104	103	97	86	11
m(GPS) = 4	112	68	44	184	104	80	110	87	23
m(GPS) = 5	100	68	32	183	105	78	101	87	14

that do not satisfy the common support restriction and therefore are removed from the analysis. At the final step, we retain only observations i that survive the common support filtering in all treatment groups. In Table 5 we report the number of observations retained in each group that satisfy the common support condition. Taken together, only 1296 out of 1659 observations can be considered as comparable in terms of their characteristics and hence are retained for a further analysis.

In order to check whether the balancing property of the constructed GPS can be warranted in our data, we subdivide each group into blocks of approximately the same size corresponding to quintiles of the respective GPS. The resulting cell sizes of each block are reported in Table 6. The testing procedure of the differences in means for each variables and for each treatment group conditioning on the GPS values is conducted in the following two steps. In the first step, five tests for the differences in means are conducted for each block. Then in the second step the computed block-specific differences in means are combined using the total number of observations in each block as weights. The balancing properties of covariates adjusted for the GPS are reported in Table 7. Compared to the results for unadjusted covariates reported in Table 4, a substantial improvement can be observed, as only three test statistics exceed the nominal 5% significance level. Hence, we can conclude that the generalised propensity score is appropriately defined.

Table 7 Covariance balance, t-statistics (GPS-adjusted)

	Dose group 1	*Dose group 2*	*Dose group 3*
$\ln \text{CAPEX}$	-2.43	-0.51	2.12
$[\ln \text{CAPEX}]^2$	-1.52	-0.58	1.84
$\ln \text{MCAP}$	-0.84	0.85	0.33
$[\ln \text{MCAP}]^2$	-0.70	0.82	0.18
$\ln \text{OP}^+$	-0.92	-0.47	1.50
$[\ln \text{OP}^+]^2$	-0.98	-0.28	1.45
$\ln \text{OP}^-$	0.36	-0.87	0.06
$[\ln \text{OP}^-]^2$	0.35	-0.84	0.16

References

Antonucci, T., & Pianta, M. (2002). Employment effects of product and process innovation in Europe. *International Review of Applied Economics, 16*(3), 295–307.

Arellano, M., & Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies, 58*(2), 277–297.

Becker, S. O., Egger, P. H., & von Ehrlich, M. (2012). Too much of a good thing? On the growth effects of the EU's regional policy. *European Economic Review, 56*(4), 648–668.

Bia, M., Flores, C. A., & Mattei, A. (2011). *Nonparametric estimators of dose-response functions*. LISER Working Paper Series 2011-40. Luxembourg Institute of Socio-Economic Research.

Blundell, R., & Bond, S. (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics, 87*(1), 115–143.

Bogliacino, F. (2014). Innovation and employment: A firm level analysis with European R&D Scoreboard data. *Economia, 15*(2), 141–154.

Bogliacino, F., Piva, M., & Vivarelli, M. (2012). R&D and employment: An application of the LSDVC estimator using European microdata. *Economics Letters, 116*(1), 56–59.

Bogliacino, F., & Vivarelli, M. (2012). The job creation effect of R&D expenditures. *Australian Economic Papers, 51*(2), 96–113.

Brandsma, A., & Kancs, D. (2016). RHOMOLO: A dynamic general equilibrium modelling approach to the evaluation of the EU's R&D policies. *Regional Studies, 49*, 1340–1359.

Bruno, G. S. F. (2005). Estimation and inference in dynamic unbalanced panel-data models with a small number of individuals. *Stata Journal, 5*(4), 473–500.

Bun, M. J. G., & Kiviet, J. F. (2003). On the diminishing returns of higher-order terms in asymptotic expansions of bias. *Economics Letters, 79*(2), 145–152.
Ciriaci, D., Moncada-Paterno-Castello, P., & Voigt, P. (2013). Innovation and job creation: A sustainable relation? Working Papers on Corporate R&D and Innovation 2013-01, European Commission, DG Joint Research Centre.
Egger, P. H., von Ehrlich, M., & Nelson, D. R. (2012). Migration and trade. *The World Economy, 35*(2), 216–241.
Hall, B. H., Lotti, F., & Mairesse, J. (2008). Employment, innovation, and productivity: evidence from Italian microdata. *Industrial and Corporate Change, 17*(4), 813–839.
Heinrich, C., Maffioli, A., & Vázquez, G. (2010). Impact-evaluation guideline. Technical Notes No. IDB-TN-161, Inter-American Development Bank.
Hirano, K., & Imbens, G. W. (2004). The propensity score with continuous treatments. In A. Gelman & X.-L. Meng (Eds.), *Applied Bayesian modeling and causal inference from incomplete-data perspectives* (pp. 73–84). Chichester: Wiley.
Ivanova, O., Kancs, D., & Thissen, M. (2019). EU Economic Modelling System: Assessment of the European Institute of Innovation and Technology (EIT) Investments in Innovation and Human Capital, Publications Office of the European Union, Luxembourg. https://doi.org/10.2791/184008
Kancs, D., & Ciaian, P. (2011). Modelling the flow of knowledge and human capital: A framework of innovative capital. *International Journal of Public Policy, 7*(1), 134–160.
Kancs, D., & Siliverstovs, B. (2016). R&D and non-linear productivity growth. *Research Policy, 45*, 634–646.
Kancs, D., & Siliverstovs, B. (2019). Employment Effect of Innovation. JRC Working Papers in Economics and Finance 2019/2, Luxembourg: Publications Office of the European Union. https://doi.org/10.2760/52757. JRC114588
Lachenmaier, S., & Rottmann, H. (2007). Employment effects of innovation at the firm level. *Journal of Economics and Statistics, 227*(3), 254–272.
Marx, K. (1867). *Das Kapital, Kritik der politischen Oekonomie*. Hamburg: Erster Band.
Moncada-Paterno-Castello, P., Ciupagea, C., Smith, K., Tubke, A., & Tubbs, M. (2010). Does Europe perform too little corporate R&D? A comparison of EU and non-EU corporate R&D performance. *Research Policy, 39*(4), 523–536.
Pianta, M. (2004). Innovation and employment. In J. Fagerberg, D. C. Mowery, & R. R. Nelson (Eds.), *Handbook of innovation* (Chapter 22, pp. 568–598). Oxford University Press.
R Development Core Team. (2008). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.
Rosenbaum, P. R., & Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika, 70*(1), 41–50.

Smolny, W. (1998). Innovations, prices and employment: A theoretical model. *Journal of Industrial Economics, 46*(3), 359–381.

Van Reenen, J. (1997). Employment and Technological Innovation: Evidence from U.K. Manufacturing Firms. *Journal of Labor Economics, 15*(2), 255–84.

Young, A. (1993). Substitution and complementarity in endogenous innovation. *Quarterly Journal of Economics, 108*(3), 775–807.