Radical innovation and firm productivity growth

Andrea Fronzetti Colladon[*] Barbara Guardabascio[†]
Ludovica Segneri[‡] Alessandro Sterlacchini[§]
Francesco Venturini[¶]

January 21, 2026

Abstract

We investigate the dynamic effects of radical innovation on firm productivity growth and its main transmission channels. Using panel data on over 28,500 Italian firms from 2012 to 2020, and analysing the text of their patents, we develop a measure of radical innovation based on the novelty content of firm innovation with respect to the state of technological knowledge. We examine how firm productivity growth responds to the launch of a novelty, finding out that output per worker growth initially slows down but accelerates in the following years. The cumulative effect of radical innovation over time is positive and economically significant: firms that introduce novel technologies experience, on average, an 8% faster rate of labour productivity growth than those without such innovations. We document that the effect of radical innovation is causal through a Difference-in-Difference regression that considers the arrival of novelties as an event of technological discontinuity and, alternatively, using an instrumental variable regressions in which, for identification, we exploit variation in firm access to local public knowledge. Finally, we illustrate that the impact of novelties on labour productivity growth is primarily driven by employment adjustments and by the performance of small firms.

Keywords: Novelty; Patent text analysis; Labour productivity growth; Difference-in-Difference

JEL: D24; O32; O34

[*]Roma Tre University (IT).
[†]University of Perugia (IT).
[‡]University of Urbino Carlo Bo (IT).
[§]Politechnic University of Marche (IT).
[¶]Corresponding author: francesco.venturini@uniurb.it. University of Urbino Carlo Bo. Via Saffi 42 Urbino (IT). National Institute of Economic and Social Research (NIESR, UK). Economic Statistics Center of Excellence (ESCoE, UK). Centre for Innovation Research, Lund University (CIRCLE, SE). We acknowledge financial support under the National Recovery and Resilience Plan (NRRP), Mission 4, Component 2, Investment 1.1, Call for tender No. 1409 published on 14.9.2022 by the Italian Ministry of University and Research (MUR), funded by the European Union – NextGenerationEU– "Digital INNOvation: an assessment of the relationship between DATA and other intangible assets and productivity (INNODATA)" – CUP J53D23006900006.

1 Introduction

Firm innovation is primarily driven by profit-seeking motives and, to this end, companies strive to develop new products, production processes, or business models. During periods of major technological breakthroughs or economic turbulence, firms engage in research projects that may result in the creation of innovations that fundamentally alter their competitive advantage and change the trajectory of technological practice (Schumpeter, 1934, Nelson and Winter, 1982, Dewar and Dutton, 1986). Radical innovations are novel as they depart from the existing technological base, but are also impactful as they affect future technological development by shaping how subsequent innovators design and produce new technologies (Dahlin and Behrens, 2005, Cohen, 2010, Archibugi et al., 2013). Radical innovations rely on previously unknown knowledge or on original re-combinations of existing knowledge (Usher, 1954, Fleming, 2001). Such knowledge may derive from several domains and can be developed both inside and outside the technological boundaries of the firm (Schoenmakers and Duysters, 2010, Castaldi et al., 2015).

The literature on the role of radical innovation is extensive. Seminal contributions identify radical innovations as key drivers of the emergence of new technological paradigms (Dosi, 1982) or the uptake of industrial revolutions (Mokyr, 1990, Lipsey et al., 2005). Another pioneering line of research looks at the technological discontinuity between innovations (Carpenter et al., 1981) and assesses how the introduction of radical innovations affects firm performance and reshapes market structure when incumbent firms are challenged by new entrants (Henderson, 1993, Christensen and Rosenbloom, 1995, Christensen, 2013). Other studies, following Dewar and Dutton (1986) and Tushman and Anderson (1986), examine the determinants of radical and incremental innovation and analyze how their interaction shapes industry emergence and evolution. An influential stream of structuralist research focuses on novelty, conceptualizing innovation as the exploration of "adjacent possibles" that enable advances into domains that were previously inaccessible. The emergence of novelties follows statistical laws arising from the recombination of past discoveries (Kauffman, 2000; Tria et al., 2014). More recent microeconomic works investigate the challenges firms encounter in initiating radical innovation projects (Rice et al., 2001), or in commercializing radical technologies, highlighting the strategic responses developed to address these challenges, such as those stemming from customers' limited information (Nasirov and Castaldi, 2025).

Measuring radical innovation and assessing its impact on firm performance, however, remains a challenging task. On the one hand, radical innovations exhibit a degree of novelty that is difficult to capture using the standard methods employed by economists and technologists. On the other hand, evaluating the impact of radical innovations on the development of later technologies may be constrained by delayed data accessibility, thereby limiting the usefulness of such information for strategic decision-making. Accordingly, although there is considerable interest in the literature, evidence on firm-level productivity effects of radical innovations, as well as their timing, remains largely scarce. This paper aims to fill this gap and, using a large panel dataset of over 28,500 Italian companies observed between 2012 and 2020, investigates the dynamic effects of radical innovation on output per worker growth and its mechanisms of transmission. We define radical innovations as novelties, examining the firm's technological distance from the existing state of knowledge, which we measure through the analysis of patent documents. Unlike impactful innovations, which are typically identified ex post looking at their effects on future technological development, novelties are identified through an ex ante assessment. This allows to study radical innovations without requiring a long observation horizon and avoids potential success bias in the analysis. Novelties are therefore ideal for understanding the strategic usage of innovation and identifying the mechanisms driving the effects of radical innovations on firm performance (Verhoeven et al., 2016).

We find that the growth rate of output per worker initially declines following the introduction of novelties but accelerates afterwards. The cumulative (net) effect of these innovations on productivity growth is positive and economically meaningful: firms introducing novel technologies experience, on average, an 8% higher productivity growth rate than firms without such innovations. This effect appears to be primarily driven by adjustments in employment.

We show that our novelty measure conveys meaningful information and constitutes a suitable tool for analyzing the firm-level effects of radical innovation, outperforming alternative measures based on citation flows or those capturing the impactful and breakthrough nature of patented innovations. Beyond measurement issues, we document that our results are robust to a wider range of econometric challenges, such as omitted variables, reverse causality and selectivity. Specifically, to understand whether the link between radical innovation and firm performance is causal, we supplement the baseline analysis with an event study tracing how labour productivity growth changes after the launch of novelties exploiting the technological discontinuities generated by these innovations. We complement the event analysis with an instrumental variables regression, where we predict the probability of finding a novelty with the firm's potential access to local public knowledge. Further, we document that our results are not driven by selectivity by implementing a matching procedure that pairs companies that introduce novelties with companies with similar characteristic without such innovations.

We conclude the study by examining the heterogeneity in the effects of radical innovation between sectors and firm-size groups. Our findings show that the dynamic effects of radical innovations are positive and significant in the service sector, but not in the manufacturing sector. In the latter, employment expansion immediately following the introduction of radical innovations is found to dampen labour productivity growth; however, subsequently, employment dynamics does not reverse, and productivity growth remains stagnant. Most importantly, we find that the main pattern of effects of radical innovation on labour productivity growth, which are negative in the short run but positive in the medium run, is primarily driven by small firms.

Our work makes two main contributions to the literature. First, it relates to studies that investigate the characteristics of radical innovations and their capacity to shape the course of technological development (Hall and Trajtenberg, 2004; Petralia, 2020). Specifically, we connect to the growing body of works that examine scientific productivity of top researchers (Arts and Veugelers, 2020), trace breakthrough innovations (Kelly et al., 2021) and map technological interdependence (Fronzetti Colladon et al., 2025) through the analysis of patent text (see also Arts et al., 2021). By bringing this type of analysis under the lens of firm performance, we shed light on the micro-foundations of these processes. While patent text analysis has gained increasing popularity and has proven to be a valuable tool for identifying innovation characteristics, evidence remains limited on the relationship between radical innovations (novelties) and firm performance. Second, we contribute to the debate on the widening productivity gap between frontier companies and the bulk of lagging companies (Criscuolo et al., 2019). While laggards often struggle to cope with the technological progress of market leaders, a segment of highly innovative, predominantly smaller, companies may instead be able to seize opportunities arising from new technologies. By tracing the dynamics of the effect of novelties, we show that productivity gains materialize only once the input adjustment following the introduction of radical innovations is complete (Brynjolfsson et al., 2021). In this respect, we provide evidence on a potential microeconomic mechanism underlying the slowdown in aggregate productivity growth (Fernald et al., 2025).

The remainder of the paper is structured as follows. Section 2 provides a review of the literature. In Section 3, we introduce the empirical model to be tested. Section 4 describes the data collection process and the construction of our variables. We then present descriptive statistics in Section 5 and econometric results in Section 6. Concluding remarks are provided in Section 7.

2 Literature review

A vast body of literature has explored how innovation can be measured, its key determinants, and its effects on firm performance (see, among others, Castellacci and Zheng, 2010; Bontempi and Mairesse, 2015; Ugur et al., 2016). Although many studies highlight positive long-term effects and the transformative potential of innovations (Hall et al., 2010), short-term impacts are often more nuanced. Innovations, even when they are not radical, often require significant organizational adjustments, new skill acquisition, and process reconfigurations, which may temporarily disrupt productivity. Chappell and Jaffe (2018), for instance, identify a negative short-run effect, driven by implementation lags, coordination costs, and learning dynamics. These transitional frictions characterize investment in R&D and other intangible assets and may result in what has been later described as a "J-curve effect", where an initial decline in productivity is followed by sustained gains as the innovation becomes embedded in the firm's operations. The delayed effect can also depend on initial mismeasurement of output or co-investment associated with intangible assets (Brynjolfsson et al., 2021).

Reliable measurement of firm-level innovation is essential for assessing its impact. Numerous studies employ R&D expenditure as an input indicator of innovative activity, while others have considered output measures, such as patent counts or the introduction of new products and processes. Crépon et al. (1998) pioneered this stream of research by developing a structural model, building on "knowledge production function" framework (see Griliches, 1990), in which R&D positively influences patenting performance, which in turn enhances firm productivity.

An established stream of research has documented that patents, in particular those that make or receive a large number of citations, convey pivotal information for identifying the most technically and commercially valuable innovations (Carpenter et al., 1981; Trajtenberg, 1990; Albert et al., 1991). More recent studies employ patent indicators adjusted for quality and relate them to firm performance (Lanjouw and Schankerman, 2004) or market value (Hall et al., 2005). Bloom and Van Reenen (2002) show that citation-weighted patents outperform raw patent counts in explaining labour productivity differences among UK firms. Another key result is that important patents exert an immediate effect on firms' market value, while their impact on productivity emerges only over time. Market value is a forward-looking indicator that reflects investors' expectations about a firm's ability to generate future profits. It may therefore respond promptly to the granting of intellectual property rights over high-quality or high-value inventions. By contrast, productivity improvements materialize only once the invention is effectively transformed into an actual innovation, a process that typically requires additional investments in skills, capital equipment and business reconfiguration. Balasubramanian and Sivadasan (2011) find similar results for the US, highlighting the high information content of forward citations and their ability to identify high-value patents. Moreover, these authors show that firms experience an increase in size (in terms of capital stock and employment) immediately after the application for a high-value patent, while the contemporaneous effect of these innovations on productivity, although positive, is comparatively weaker.

A key feature of the literature is that patent radicalness is typically measured using backward citations. According to Ahuja and Lampert (2001), the fewer citations an innovation makes to earlier patents, the greater its technological discontinuity from prior technologies. Another influential contribution in this respect is Shane (2001), who defines radicalness as "the degree to which an invention differs from previous inventions in the field" (ibid, p. 207). This indicator is inherently backward-looking, as it refers only to prior patents, and does not indicate whether a patent is radical relative to other patents filed in the same technological class and period. Dahlin and Behrens (2005) define radicalness in terms of the novelty, uniqueness, and potential impact of patented inventions on subsequent innovations. They observe the citation structure before, during, and after a patent's filing to assess whether, relative to other patents, an invention can be classified as radical.

Despite their widespread use as a metric for patents, citation-based indicators have several well-known limitations. They are influenced by the discretion and expertise of both inventors and patent examiners, with the latter being responsible for adding a substantial share of citations. For instance, applicants may omit certain citations for strategic reasons, while examiners may not always identify or include all relevant prior art (cf. Funk and Owen-Smith, 2017).[1] Moreover, forward citations are characterized by discrete, discontinuous, and highly skewed distributions: a small number of patents receive numerous citations, whereas most receive few or no citations. However, this does not necessarily imply that uncited patents are unimportant to trace radical inventions. Last but not least, forward citation data are not immediately available, as they can be observed only several years after a patent's publication (Ashtor, 2019). This poses a significant limitation for firm-level analyses, where the time window of observation is also restricted by the limited availability of complementary firm-level information such as, for instance, balance sheets, tax records, etc.

Most of these limitations are overcome by indicators based on the textual analysis of patent documents. Ashtor (2019) applies latent semantic analysis to patent claims to measure similarity among inventions within the same technology–time cohort, where lower similarity indicates a greater degree of radicalness. This measure is found to be correlated with forward citations and market value at firm level in the US. Kelly et al. (2021) introduce a new measure of invention importance, defined as breakthrough, to study technological development in the United States. A patented invention is considered as breakthrough if its content is distinct from earlier patents (i.e. novel) but is similar to the subsequent ones (i.e., impactful). Fronzetti Colladon et al. (2025) measure interdependence among technology sectors using patent text similarity and compare it with similarity based on patent citation flows. While citation-based similarity is found to capture common core technologies, text-based similarity reflects a broader range of technological characteristics of innovation. A methodological review of the first generation of patent quality indicators is provided in Squicciarini et al. (2013), while Arts, Hou, and Gomez (Arts et al.) offers an overview of more recent methods and quality indicators derived from the textual analysis of patent documents. Further methodological details on the use of patent text analysis to measure innovation characteristics are provided in Section 4.

[1] To circumvent these limitations, Verluise et al. (2025) examine in-text citations and find that they exhibit greater semantic proximity than front-page citations, likely because they more accurately reflect inventors' knowledge.

3 Econometric model

The primary goal of this paper is to understand whether the rate of productivity growth accelerates or slows down in the aftermath of the launch of radical innovation and how this effect changes over time. We explore this issue using panel data for Italian firms spanning from 2012 to 2020 and estimate a Difference-in-Difference (DiD) regression shaped as follows (lower cases denote logged variables):

$$\Delta\, lp_{it} = \alpha_i + \beta N_{it} + \gamma x_{it} + \mu_{jt} + \epsilon_{it}, \tag{1}$$

where i denotes firms and t years. The dependent variable is the annual log change in labour productivity, whilst N is a time-varying proxy for radical innovation. In line with the existing literature, we adopt an absorbing treatment framework and assume that once a firm introduces a novelty, it remains treated for the remainder of the observation period (Athey and Imbens, 2022). N is therefore a treatment dummy equal to 1 from the year of introduction of a radical innovation onward and 0 otherwise. Since companies introduce novelties in different points in time, our regression model is staggered. β is the average post-treatment effect of radical innovation, namely the absolute change in labour productivity growth (expressed in decimals). To exclude omitted variables' problems, our baseline model includes a set of standard controls (all in logs): (i) the lagged level of labour productivity to capture the effect of convergence (or divergence) process; (ii) the capital-labour ratio to filter the effect of capital intensity of production; and (iii) the company's patent portfolio used as a proxy for technological capabilities, defined as number of patents per worker. To capture unobserved heterogeneity, we include firm fixed effects (α_i) and sector-year fixed effects (μ_{jt}). The former measure the impact of those firm characteristics that are unobservable and do not change over time (patent propensity, etc.). The latter isolate the effect of sector-specific factors, such as technological shocks, demand shifts and, particularly important in our case, the dynamic in industry-level prices as balance sheets data are expressed in current prices. ϵ_{it} is the classical spherical error term.

In the analysis above, the treatment is considered exogenous with respect to the outcome variable. However, one cannot exclude that firms engage in innovation processes as a result of earlier productivity gains or, alternatively, to escape from a low productivity growth regime. Reverse causality would make the coefficient of the treatment variable biased. To overcome this risk, we estimate our DiD specification (Eq. (1)) using two further regression methods. As a first approach to address causality, we use the Local-Projections DiD regression (LP-DiD, Dube et al., 2025). This amounts to estimating a series of forward effect regressions assessing how productivity growth evolves around the innovative event, defined as the year of the launch of radical innovation. Specifically, for each horizon h, we regress the change in productivity growth between t and $t+h$ on treatment status at time t:

$$\Delta\, lp_{it+h} - \Delta\, lp_{it} = \sum_{h=-2}^{2} \beta_h N_{it} + \gamma x_{it} + \sum_{p=-1}^{1} \zeta_p\, N_{it-p} + \mu_{jt} + \varepsilon_{it}. \tag{2}$$

in which β_h denotes the single-year Average effect of Treatment on the Treated (ATT) at horizon h. We present the results of the LP-DiD regression as an event analysis, comparing the change in productivity growth between treated and non-treated firms, with respect to the pre-treatment period. The control group consists of firms that had not yet introduced, or had never introduced, any novelties during the period of analysis. Along with confounders (x's), Eq. (2) includes the one-year lag of the dependent variable to exclude selection in the treatment as firms with given productivity performance may have more (or fewer) incentives to innovate and hence have a higher (or lower) probability to develop a novelty. Eq. (2) also accounts for dynamics by admitting one-year lags and leads of the treatment variable ($p = 1$). Lags neutralize the bias potentially associated with the delayed response of productivity growth to novelties. Leads exclude anticipation effects. For instance, companies may change in advance their organisational setting and managerial practices expecting the arrival of novelties. These changes may, in turn, increase the rate of productivity growth. Since we control for both cross-firm differences in initial productivity levels and their dynamics over time, the introduction of a novelty can be interpreted as a discontinuity in the firm's trajectory with respect to the overall state of technological knowledge. This constitutes the identifying assumption underlying our LP-DiD regression and is consistent with the approach adopted by Nilsen and Raknerud (2024).

As a second approach used to address causality, we run a two-stage instrumental variables regression. In the first stage, we predict variation across companies in the development of novelties using information on firm's potential exposure to local public knowledge. A well-established body of literature shows that firms exhibit greater innovative capacity when they interact with public research institutions, as such interactions allow access to basic knowledge and research in emerging technological domains that firms are unable to develop internally. Arts and Veugelers (2020) document that business-sector scientists collaborating with academic researchers, proxied by co-authored scientific publications, develop a greater number of novel (and impactful) patented innovations. Gómez et al. (2020) find that research conducted jointly with public research institutions leads to develop high-novelty innovations, which are often capable of capturing larger market shares. We measure access to local public knowledge using two distinct region-level variables: the percentage share of public R&D on GDP (below denoted as *prd*) and the logged number of scientific publications co-authored by business and university researchers (denoted as *copub*). Notably, the former is an input measure of the public research effort; the latter is a proxy for the scientific output jointly achieved by companies and higher education institutions:

$$N_{it} = \alpha + prd_{rt} + copub_{rt} + \vartheta z_{it} + \mu_j + \mu_{jt} + \mu_{rt} + \varepsilon_{it} \tag{3}$$

$$\Delta\, lp_{it} = \alpha_i + \sum_{h=-2}^{2} \beta_h \hat{N}_{it} + \gamma x_{it} + \mu_{jt} + \epsilon_{it}. \tag{4}$$

The innovation equation, Eq. (3), is estimated with a pooled probit regression with industry fixed effects, μ_j, industry-by-year dummies, μ_{jt}, and region-by-year dummies, μ_{rt}. μ_j helps capture sectoral differences in patenting (or R&D) propensity, μ_{jt} collects common changes at industry level in technological opportunities, while μ_{rt} neutralize the effect of unobservable time-varying regional characteristics (public expenses, localised technological or trade shocks) affecting the firm's propensity to develop novelties. Note that using variation in local public knowledge as a predictor of novelties, rather than variation in firms' collaborations with public research institutions (co-patenting), mitigates selectivity problems in the probit regression. Moreover, to avoid omitted variables' bias, Eq. (3) includes various control variables (z_{it}), such as firm's patent and physical capital intensities, lagged productivity, along with proxies for localised competition effects or knowledge spillovers from the business sector.

The predicted value of radical innovation, $\hat{N}$, from Eq. (3) is then used in the second-stage regression as a determinant of labour productivity growth (Eq. (4)). This equation is similar to our baseline model (Eq. (2)) but is estimated as an event regression by including lagged, contemporaneous, and lead values of the treatment variable ($h = -2, ..., 2$). β_h's are the *predicted* pre-treatment and post-treatment periods' effect of novelties on labour productivity growth, that we illustrate below in a event-analysis plot. As $\hat{N}$ is a generated (non-random) variable, standard errors associated with the instrumented variable are computed with a wild bootstrap procedure based on 1,000 replications. In both Eqs. (2) and (4), insignificance of pre-treatment coefficient ($h = -2, -1$) ensures that the parallel trend assumption is satisfied. The significance of the post-treatment parameters ($h = 1, 2$) will instead prove that there is a differential productivity growth performance between companies with radical innovations (treated units) and the rest of the sample (control units).

To further mitigate selectivity problems, we also implement a matching procedure that pairs treated firms with companies with similar characteristics drawn from the group of control firms without novelties. The risk is of comparing firms that differ for reasons unrelated to radical innovation. This concern arises as we can only observe firms' productivity growth after the introduction of radical innovations, while the counterfactual productivity performance in absence of these innovation remains unobserved. As in Marioni et al. (2024), we apply propensity score matching by estimating the probit model in Eq. (3) separately for each year in the sample period. Based on the resulting propensity scores, each treated firm is matched with either five or ten nearest neighbours on the same support. We then replicate the event-study analysis in Eq. (2) on the resulting restricted sample of matched group of treated and control firms.

4 Data sources and variable definitions

The empirical analysis is conducted on a panel of 28,611 Italian firms operating across 29 industrial sectors (NACE 2-digit classification), observed over the period from 2012 to 2020. We use the Moody's Orbis Intellectual Property (IP) database, which merges firm-level balance sheet information with patent data originally drawn from the World Patstat database. The selection of the sample is based on two sets of criteria. On the financial side, we retain firms with available data on value added, employment and total fixed assets. See Section 4.1 for details. On the innovation side, we consider only patents with an English-language text (title and abstract) and for which the current patent owner coincides with the original applicant, thereby ensuring unambiguous attribution of inventive activity. The final number of patent documents included in the text analysis is 29,332. For each patent, we construct two novelty indicators. The first indicator measures the textual distance between a patent and the technological content of patents filed by firms in the sample over the previous five years. The second indicator, which is used for comparison purposes, relies on information on backward citations.[2] Details on the construction of the novelty indicators are provided in Section 4.2. We define as treated all firms that hold at least one patent with a radical innovation indicator above the 90th percentile of the patent novelty distribution. Overall, 1,510 firms satisfy these conditions when novelty is measured using the text-based indicator and are therefore classified as treated, while 2,377 firms are considered as treated when novelty is measured using the backward citation–based indicator. All remaining firms, i.e., those that fail to meet at least one of the above condition, are assigned to the control group (non-treated). As discussed below, we assess the validity of our sample selection through a matching procedure.

[2]Backward Citations are extracted from Google Patents Public Datasets through Google BigQuery API with the google.cloud.bigquery and pandas Python libraries. In some cases, metadata is complemented using the "requests" and "BeautifulSoup" Python packages for direct parsing of Google Patents web pages.

4.1 Balance sheet data

For each firm, we construct the following variables. Labour productivity is measured as value added per employee and total sales per employee. Capital stock per worker is estimated annually using the perpetual inventory method, combining total fixed-asset investment with annual capital depreciation derived from the balance sheets, and is then scaled by firm-level employment. We define the cash-flow ratio as cash and cash equivalents over output; the intangible ratio is defined as intangible fixed assets over total fixed assets, and the debt ratio as current liabilities over total assets. Patents per worker are constructed by counting, for each firm-year, the number of patents recorded in Orbis IP divided by employment. We also include two time-varying, context-level indicators reflecting competition pressure and the breadth of knowledge spillovers. Competition effects are proxied by the number of innovative firms operating in the same region–sector pair as the focal firm. Knowledge spillovers are proxied by the total number of patents filed in that same region–sector pair. Both variables are constructed excluding the focal firm.

4.2 Patent data

We construct two measures of novelty for radical innovation. The first measure is derived using text analysis techniques (NLP) applied to the titles and abstracts of patents. These condense the most technically relevant part of patent documents, being largely homogeneous over time and across jurisdictions (Fronzetti Colladon et al., 2025). The patent text–based measure is hereafter referred to as N_{NLP}. The second measure is derived from backward citations and is labelled hereafter as N_{BWC}.

N_{NLP} is obtained by comparing the title and abstract of each of the 28,611 patents filed between 2012 and 2020 by our sample of firms, with the bulk of patent applications filed in the five-year interval preceding their priority date. For this purpose, we extract a total of 51,705 patent texts with priority dates between 2007 and 2020 from Orbis IP, ensuring that every patent in the period 2012–2020 has a complete five-year backward window, as required to construct our novelty indicator. For each patent, we transform the patent texts into numerical vectors, where each dimension reflects the frequency of a given term (Term Frequency, TF) and the importance of that term is computed considering its occurrences in earlier patents (Past Inverse Document Frequency, PastIDF). The term-weighted representation is defined as:

$$\text{TFIDF}_{p,w} = \text{TF}_{p,w} \times \text{PastIDF}_w \tag{5}$$

in which $\text{TF}_{p,w}$ measures how often a term w appears in a given patent p. PastIDF_w captures the rarity, and thus the distinctiveness, of the term w across all patents filed in the preceding five years, based on its document frequency. Common terms receive low weights, as they contribute little to novelty. Formally:

$$\text{PastIDF}_w = \log\left(\frac{P+1}{1+d(w)}\right) + 1 \tag{6}$$

where P is the number of patents filed in the five years before patent p, and $d(w)$ is the number of those patents containing term w. By construction, the weighting scheme reflects only the technological vocabulary available prior to the focal patent p, avoiding distortions associated with the wording of contemporaneous filings. In this way, we represent each patent as a numerical vector, which is instrumental in the later calculation of cosine similarity between documents. We also remove overly common terms that appear in 85% or more of documents. L2 normalization is then applied to ensure comparability across documents of different lengths. Our novelty index, N_{NLP}, for patent p is defined as the maximum textual distance from prior patents:

$$\mathrm{N}_{\mathrm{NLP}}(p) = 1 - \max_{q \in \{P\}} \rho_{p,q} \tag{7}$$

where P is the bulk of patents filed in the five years preceding p, and $\rho_{p,q}$ is the cosine similarity between patents p and q. A high value for N_{NLP}, indicates that the focal patent p diverges substantially from prior art, reflecting a high degree of technical (semantic) novelty. Conversely, a low score suggests proximity to at least one earlier patent, indicating that the focal patent p conveys incremental innovation. Our choice to consider the maximum distance from prior art aligns with previous research (Gerken and Moehrle, 2012) and adds to other approaches that measure the average distance of the focal patent from past and future innovations (Arts et al., 2021, Kelly et al., 2021). These might inflate the degree of patent novelty for innovations lying in small or niche domains: in these areas, a focal patent will, on average, lie far from the large out-of-field corpus, making even incremental advances appear unusually "distant". In contrast, N_{NLP}, constructed as maximum distance from earlier patent texts, isolates the most substantial deviations from prior knowledge, thereby effectively detecting patents that embody genuinely novel methodological advances.

To identify radical innovations, we examine the distribution of N_{NLP} scores across all patents in each year of our sample and define as novel those patents that fall above the 90th percentile. The firms owning these patents are classified as treated firms. As shown by Verhoeven et al. (2016), the distribution of technological novelty is highly skewed, with a small subset of patents exhibiting disproportionately high radicalness. Selecting the 90th percentile thus offers a pragmatic trade-off: it is sufficiently restrictive to capture genuinely disruptive innovations, while ensuring that the results are not driven by a handful of firms. This approach is consistent with established empirical practices in innovation research (e.g., Kelly et al., 2021). Robustness checks assessing the sensitivity of our estimates to the chosen threshold are reported in Section 6.

For comparison purposes, we construct a measure of novelty based on backward citations, BWC, defined as follows:

$$\mathrm{N}_{\mathrm{BWC},p} = 1 - \frac{\mathrm{BWC}_p}{\max\left(\mathrm{BWC}_t\right)} \tag{8}$$

where BWC refers to the sum of the backward citations included in the focal patent p, t represents the filing year of patent p, while $\max(\mathrm{BWC}_t)$ is the maximum number of backward citations observed among all patents filed in the same year (cohort) as the patent p. As in Ahuja and Lampert (2001) and Banerjee and Cole (2010), the underlying idea is that a high number of backward citations may reflect a strong reliance on existing knowledge. In contrast, a low number (i.e., a higher value of N_{BWC}) indicates a greater degree of novelty (or originality), as the patent exhibits a lower dependence on prior technological knowledge (see also Lanjouw and Schankerman, 2001). In particular, if a patent has the highest number of citations made to earlier patents in its cohort, i.e., $\mathrm{BWC}_p = \max\left(\mathrm{BWC}_t\right)$, the indicator equals zero, signalling a low degree of novelty. Conversely, a value close to 1 indicates that the patent has made few citations, implying it has a high degree of novelty. As for $\mathrm{N}_{\mathrm{NLP}}$, we consider as novelties (radical innovations) those lying above the 90th percentile of the distribution of $\mathrm{N}_{\mathrm{BWC}}$.

In addition to the aforementioned indicators of novelty, N_{NLP} and N_{BWC}, we compute an array of text-based metrics to illustrate the distinctive information content conveyed by our main measure of radical innovation. First, we consider the average mean value of firm novelty from earlier technologies, N_{NLP}^{Mean}. This is built following the same procedure as N_{NLP}, but it uses the average distance of the focal patent from those filed in the previous five years, rather than considering the maximum distance. Second, we construct an index reflecting the patent impact on the development of subsequent technologies, I_{NLP}. This is calculated as the average similarity of the focal patent to all patents filed in the five years following its priority date.[3] Finally, in order to identify patents that are both highly novel and highly impactful, we construct a measure of breakthrough innovation as in Kelly et al. (2021):

$$B_{\mathrm{NLP}} = \frac{I_{\mathrm{NLP}}}{1 - N_{\mathrm{NLP}}} \tag{9}$$

Technically, this measure helps distinguish patents whose terminology substantially departs from prior knowledge and is extensively reused in subsequent innovations. Economically, it quantifies the extent to which a firm's innovations differ from previous advances and influence future technological developments. The former characteristic (novelty) indicates how much a firm's innovation diverges from earlier technologies, while the latter (impact) reflects how much subsequent technologies differ from the firm's innovation. In the analysis that follows, we demonstrate that novelty indicators are particularly well-suited to capturing the effects of radical innovation at the firm level, especially when the post-innovation time horizon for assessing their impact is limited.

5 Descriptive analysis

We begin by reporting descriptive statistics in Table 1 that compare firms introducing radical innovations, as measured by the text-based novelty indicator, with those that do not (the control group), to gain initial insight into their characteristics and performance.

Table 1: **Descriptive Statistics: Text-based novelty**

	Treated (1,510 firms)			Control (27,101 firms)		
	Mean	Median	SD	Mean	Median	SD
Labour productivity growth (%)	0.28	-0.005	78.98	-0.71	-0.09	88.36
Controls						
Capital per worker (mln)	32.85	9.20	278.44	65.35	6.35	1411.17
Patents per worker	0.18	0.03	1.34	0.39	0.06	14.65
# of innovators by sector-region	10.55	0	27.63	0.93	0	8.8
# of patents by sector-region	213.69	56	326.77	142.92	27	262.8
Cash flow ratio	2.29	0.07	79.74	23.91	0.05	3974.98
Intangible ratio	0.24	0.11	0.29	0.18	0.05	0.27
Debt Ratio	0.68	0.71	0.21	0.71	0.76	0.23
Employment growth (%)	0.18	0	41.35	0.14	0	46.36
Value Added growth (%)	0.46	-0.11	65.85	0.07	-0.14	74.12

Firms that introduced radical innovations (1,510 in total) during the observation period (2012–2020) exhibit distinct structural and financial characteristics compared to the control group. On average,

[3] We construct a TF–IDF matrix that includes the focal patent and those filed within the following five years; however, the IDF values are computed using only patents from the focal year and the five preceding years. This weighting scheme ensures that newly introduced terms in the focal patent, i.e., those that later diffuse and are reused, are not underweighted, thereby capturing the patent's innovative impact more accurately.

treated firms show higher labour productivity growth (0.28% versus −0.71%). However, large standard deviations indicate highly skewed distributions, suggesting that while many firms experience little or no change in productivity, a smaller subset undergoes substantial gains or losses. This pattern suggests that only a fraction of firms experience significant productivity improvements following the introduction of radical innovations, consistent with the uneven distribution of innovation outcomes documented in the literature (Verhoeven et al., 2016). Control firms display, on average, higher levels of both capital and patents per worker than treated firms. Capital intensity amounts to 65.35 million per worker in the control group compared to 32.85 million among treated firms. A similar pattern emerges for patents per worker, with higher means in controls (0.39 versus 0.18). Our proxies for localised innovation activities reveal that radical innovators are more often located in areas with a greater density of innovative firms and innovation output, consistent with the literature on geographical clustering and localised knowledge spillovers (Jaffe et al., 1993; Audretsch and Feldman, 1996). Also, treated firms present, on average, a higher share of intangible assets (0.24 versus 0.18) and a lower debt ratio (0.68 vs. 0.71) compared to the control group. This result suggests that radical innovators are more intangible-capital intensive while relying less on debt financing, possibly indicating more cautious financial structures or stronger internal funding capacity for innovation-related investments. At the same time, financial indicators reveal substantial heterogeneity: the cash-flow ratio exhibits extreme variation across both treated and control firms, suggesting that financial conditions differ widely regardless of innovation status. This echoes previous evidence on the sensitivity of innovation to internal financial resources and capital structure (Brown et al., 2009; Hall and Lerner, 2010).[4]

6 Econometric results

This section presents estimates on the relationship between radical innovation and productivity growth. First, we assess the short-run (contemporaneous) impact of radical innovation. We then examine the dynamic adjustments of these effects by quantifying the cumulative impact of our explanatory variable over time. In addition, we validate the information content of our novelty indicator with respect to other measures of radical innovation and show the robustness of our results to the assumptions made in data construction. Not less importantly, we address potential endogeneity issues that may affect the causal interpretation of the relationship under investigation. Next, we investigate the transmission channels through which radical innovation affects labour productivity growth, specifically whether the impact operates through changes in output growth or through employment dynamics. As a final step, we examine how the impact of radical innovation varies across sectors and firm sizes, aiming to identify which types of firms benefit the most and which are less affected from novelties.

Contemporaneous effects

Table 2 presents estimates for the contemporaneous impact of radical innovation on firm productivity growth comparing results based on our text-based novelty indicator (reported in the left-hand columns) with those obtained using the backward citation-based measure (reported in the right-hand columns). In each group of estimates, we consider output growth defined in terms of sales and value added. Next,

[4]Table A.1 in Appendix presents descriptive statistics for the 2,377 treated firms and the 26,234 control firms when treatment is defined using backward citations-based novelty.

we extend the model by including controls for localised competition effects and knowledge spillovers. Finally, we expand the set of firm-level controls to include the cash-flow ratio, intangible capital intensity, and short-term debt ratio. The main drawback of incorporating the latter set of variables is that they are available for a smaller subset of firms, thereby reducing the sample size and limiting the comparability of estimates.

Table 2: **Contemporaneous effects of novelties**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Text				**Backward cites**			
Novelty (post × treatment)	-0.056***	-0.061***	-0.066***	-0.049***	-0.031	-0.058***	-0.063***	-0.044**
	(0.018)	(0.019)	(0.019)	(0.019)	(0.022)	(0.022)	(0.022)	(0.021)
Lagged productivity (log)	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***	-0.010***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Capital per worker (log)	0.132***	0.129***	0.129***	0.165***	0.132***	0.129***	0.129***	0.165***
	(0.005)	(0.004)	(0.004)	(0.006)	(0.005)	(0.004)	(0.004)	(0.006)
Patents per worker (log)	0.018***	0.020***	0.020***	0.021***	0.018***	0.020***	0.020***	0.021***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
# of innovators by region-sector (log)			0.000***				0.000***	
			(0.000)				(0.000)	
# of patents by region-sector (log)			-0.000				-0.000***	
			(0.000)				(0.000)	
Cash flow ratio				-0.081***				-0.081***
				(0.001)				(0.001)
Intangible ratio				-0.008***				-0.081***
				(0.002)				(0.002)
Debt ratio				0.165***				0.165***
				(0.012)				(0.012)
Labour productivity measure	Sales p.w.	Value added p.w.	Value added p.w.	Value added p.w.	Sales p.w.	Value added p.w.	Value added p.w.	Value added p.w.
Observations	238,873	189,802	189,802	147,338	238,873	189,802	189,802	147,338
Firms	31,431	28,611	28,611	24,229	31,431	31,431	28,611	24,229
R-squared	0.564	0.560	0.560	0.610	0.564	0.560	0.560	0.610

Notes: The dependent variable is the annual rate of change of labour productivity. Parameter estimates are obtained from OLS regressions including firm and industry-by-year fixed effects. Standard errors, reported in parentheses, are robust to heteroskedasticity and autoserial correlation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Our findings suggest that the introduction of novelties is significantly associated with a reduction in the annual growth rate of labour productivity. In terms of magnitude, the coefficients of radical innovation are economically important. At the introduction year, radical innovations can be linked with a drop between 4.4 and 6.6 percent in labour productivity growth rate in innovating firms compared to the control group. This aligns with the idea that integrating innovations with a high novelty content often requires substantial adjustments at the firm level, including organizational restructuring, staff training, or assimilation of new technologies, which may temporarily outweigh productivity gains. Table 2 offers some additional interesting insights. First, the negative contemporaneous effect of radical innovation emerges regardless of whether novelty is measured using the semantic (dis)similarity of patent documents or the more traditional measure based on backward citations. Second, the effect of novelty is similar whether output is measured in terms of sales or value added. In the following, though, we adopt the latter as our preferred measure of output as material costs cannot be consistently controlled for across a substantial number of firms. Third, the estimated effect of novelties remains robust even after controlling for both firm structural characteristics and contextual factors. Initial labour productivity is negatively associated with subsequent productivity growth, supporting a convergence pattern, whereby less productive firms tend to grow faster. Firms with higher intensities of physical or knowledge-based capital follow a faster productivity trajectory. The detrimental contemporaneous impact of radical innovation does not weaken when we control for external factors, such as localized competition effects and knowledge spillovers from the business sector. Finally, the main pattern of results remains unchanged even when including a broader set of control variables, namely intangibles, cash flow, and debt ratios.[5]

Dynamic effects

To assess how the productivity growth effects of radical innovation evolve over time, we extend the baseline specification by including up to three lags of our key regressors (Table 3). The table presents the estimated coefficients for each lag, along with the cumulative effect of radical innovation and the corresponding F-test of joint significance. As above, we alternatively define output either as sales or value added, and measure novelty using either text-based or backward citation–based indicators. The regressions also include our standard set of controls, namely, lagged labour productivity, the capital–labour ratio, and the number of patents per worker. However, for sake of brevity, the parameters associated with these variables are omitted from the table.

Table 3: **Dynamic effects of novelties**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Text					**Backward cites**				
Novelty (post × treatment)	-0.061***	-0.329***	-0.333***	-0.294***	-0.433***	-0.058***	-0.305***	-0.295***	-0.288***	-0.386***
	(0.019)	(0.031)	(0.036)	(0.043)	(0.043)	(0.022)	(0.033)	(0.038)	(0.045)	(0.046)
Novelty t-1		0.350***	0.304***	0.294***	0.486***		0.322***	0.254***	0.262***	0.361***
		(0.030)	(0.039)	(0.044)	(0.047)		(0.031)	(0.041)	(0.046)	(0.047)
Novelty t-2			0.081***	0.022	0.032			0.085***	0.062*	0.096***
			(0.023)	(0.031)	(0.030)			(0.026)	(0.036)	(0.033)
Novelty t-3				0.055**	0.058**				0.007	-0.016
				(0.024)	(0.023)				(0.027)	(0.025)
Novelty (cumulative)				0.076**	0.141***				0.044	0.055
				(0.034)	(0.030)				(0.039)	(0.035)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Labour productivity measure	Value added p.w.	Value added p.w.	Value added p.w.	Value added p.w.	Sales p.w.	Value added p.w.	Value added p.w.	Value added p.w.	Value added p.w.	Sales p.w.
Observations	189,802	189,802	166,431	142,493	179,304	189,802	189,802	166,431	142,493	179,304
R-squared	0.56	0.56	0.571	0.585	0.586	0.56	0.56	0.571	0.585	0.586

Notes: The dependent variable is the annual rate of change of labour productivity. Parameter estimates are obtained from OLS regressions including firm and industry-by-year fixed effects. Controls (in logs) include lagged labour productivity, the capital–labour ratio, and patents per worker. Standard errors, reported in parentheses, are robust to heteroskedasticity and autocorrelation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

experience slower productivity growth. However, in subsequent years, they exhibit, on average, productivity growth that is 30–40 percent higher, and this positive effect persists over time. When aggregated over the entire period, the cumulative premium in productivity growth amounts to 8–14

[5] In our sample, firms with a higher intangible intensity ratio on total fixed assets, a higher cash flow and lower financial leverage, experience slower productivity growth.

percent when is measured with the text-based novelty indicator. By contrast, the dynamic effect of radical innovation, as captured by backward citations, is relatively short-lived and becomes statistically insignificant within a three-year horizon. These findings suggest that radical innovation may entail substantial short-run adjustment processes and that these benefits materialize relatively late. The resulting trajectory is consistent, even for the magnitude of the impact, with the evidence on delayed productivity effects of patenting and intangibles documented in Bloom and Van Reenen (2002) and Chappell and Jaffe (2018), respectively. The latter work, in particular, shows that radical innovations impose initial productivity penalties due to organizational frictions and learning costs, and these effects turn positive only over time. One contribution of the present work to the literature is to corroborate this type of findings using patent text-based indicators of novelty.

Measurement issues: Alternative radical innovation indicators

In Table 4, we evaluate whether the proposed indicator captures economically valuable novelties, and whether its information content differs from the one conveyed by other popular metrics of radical innovation. To address this issue, we first examine how the productivity growth effects of novelties change when applying alternative thresholds for defining an innovation as radical. For brevity, in Table 4 and in the remainder of the work, we report only the coefficients of our main explanatory variable. In columns (2) and (3), we present estimates based on samples of treated firms classified as radical innovators using respectively the 99th and 75th percentile thresholds of the distribution of patent novelties. The overall pattern of results remains largely unchanged but, as expected, the cumulative productivity growth premium is larger under a stricter definition of radical innovation and smaller when we use a less conservative threshold. Specifically, for firms in the top one percentile of the novelty distribution, the estimated marginal effect of radical innovation is more than four times that obtained in the benchmark regression reported in column (1) (0.331 versus 0.076). By contrast, when a lower threshold is applied, the estimated productivity growth premium amounts to 0.059, as the treatment group includes a broader set of firms with less novel innovative content.

Next, we compare our main results with those obtained using alternative measures of radical innovation (see Section 4). Specifically, we consider the mean value of patent novelty in column (4), the impact indicator in column (5), and the breakthrough index in column (6). The productivity growth effect of novelty, measured as the mean distance among patent texts (column (4)), is estimated to be one-fourth larger than that obtained using the maximum text dissimilarity in column (1) (0.095 versus 0.076). This suggests that our indicator of maximum text distance provides more conservative and stable estimates for the economic impact of radical innovation.[6] Estimates in the last two regressions are based on data ending in 2015 as patent text measures reflect the semantic distance of each patent from future innovations within a five-year window. In both columns (5) and (6), the sign and magnitude of the short-run effects of novelty are broadly consistent with those in the baseline regression (column (1)). In each robustness regression, there is a positive dynamic adjustment in the effect of novelties that partially offsets the initial negative short-run impact. This adjustment appears incomplete owing to the shorter time span over which the indicators of patent impact and breakthrough are computed. This makes novelty measures particularly suitable when data availability limits the time span of the analysis and, more generally, when the objective is to identify the technological sources of a firm's competitive

[6] See the Online Appendix for further results based on the novelty indicator based on mean textual distance.

advantage for strategic management decisions.

Table 4: **Robustness to radicalness thresholds and alternative indicators**

	(1)	(2)	(3)	(4)	(5)	(6)
Novelty (contemp.)	-0.294***	-0.311**	-0.320***	-0.418***	-0.201***	-0.250*
	(0.043)	(0.128)	(0.031)	(0.056)	(0.060)	(0.138)
Novelty (cumulative)	0.076**	0.331***	0.059**	0.095**	0.049	0.081
	(0.033)	(0.115)	(0.025)	(0.043)	(0.091)	(0.163)
Radical innovation indicator	Max novelty N_{NLP}	Max novelty N_{NLP}	Max novelty N_{NLP}	Mean novelty N^{Mean}_{NLP}	Impact I^{NLP}	Break-through B^{NLP}
Threshold (percentiles)	90th	99th	75th	90th	90th	90th
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	142,493	142,493	142,493	142,493	44,198	44,198
R-squared	0.585	0.585	0.586	0.585	0.755	0.755

Notes: The dependent variable is the annual rate of change of labour productivity. Parameter estimates are obtained from an OLS regression including firm and industry-by-year fixed effects. Controls (in logs): Lagged labour productivity. Capital-labour ratio. Patents per worker. Standard errors robust to heteroskedasticity and autocorrelation in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Econometric issues: Causality and selectivity

A first key issue in the analysis conducted thus far is that firms' ability to generate radical innovations is treated as exogenous with respect to their productivity outcomes. This assumption reflects the nature of radical innovations, which represent discontinuities from new technologies previously developed in the market. However, one may question about the validity of this assumption within a static regression framework, arguing that the achievement of radical innovation may be the result of prior productivity performance, or may be related to factors which are unable to account for. As a consequence, the estimated marginal impact of novelties may be upward biased because of such endogeneity issues.

To exclude this risk, we first perform an event analysis and estimate our DiD specification as Local Projections regressions, as of Eq. (2). With this approach, we still assume the arrival year of radical innovation as an event (treatment). However, we remove the possible simultaneity by using one-year lags and leads of the treatment variable, along with controlling for lagged levels of per-worker output, capital-labour ratio, and patents per worker (all in logs). This identification strategy is consistent with Nilsen and Raknerud (2024) who study the effect of first patenting along several dimensions of firm performance (productivity, output, employment, etc.). Figure 1 (top panel) shows the single-year coefficient yielded by LP-DiD regressions. Even with this approach, a clear J-shaped pattern emerges in the productivity growth effects of novelties. Before the treatment, there is no significant difference in performance between treated and untreated firms. After the treatment, the productivity growth rate of firms introducing novelties declines by more than 24%, before rebounding by approximately the same

Table 5: **Drivers of radical innovation: Probit regression**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Regional public R&D	0.881***		-0.0134		1.120*		-0-109
	(0.075)		(0.101)		(0.08)		(0.105)
Regional scientific co-authorships		0.371***		0.057***		0.471***	0.0879**
		(0.031)		(0.028)		(0.033)	(0.044)
# of patents by region-sector (log)			-0.0001	-0.0001*	-0.0001	-0.0001	-0.0002
			(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
# of innovators by region-sector (log)			0.002***	0.002***	0.002***	0.002***	0.002***
			(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-by-Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-by-Region FE	Yes	Yes	No	No	No	Yes	Yes
Observations	102,529	102,529	105,362	105,362	102,529	102,529	105,362
Pseudo R-squared	0.094	0.094	0.148	0.148	0.153	0.153	0.148
Log-likelihood	-4119	-4119	-3894	-3894	-3854	-3854	-3893

Notes: The dependent variable is a dummy variable indicating whether the company introduces a novelty in a given year. Parameter estimates are obtained from pooled probit regressions including industry, industry-by-year, or region-by-year fixed effects. Controls (in logs) include lagged labour productivity, the capital–labour ratio, and patents per worker. Standard errors are clustered at the regional (NUTS2) level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

magnitude. Overall, the post-treatment productivity growth acceleration is of 3.4% (standard error 0.018), which is significant at a 10% level.

As an alternative identification strategy, we perform the IV regression described in Section 3. In the first stage (Table 5), we run a probit model to estimate the effect of public R&D intensity and scientific publications co-authored between business-sector and academic scientists at regional level on a firm's probability of introducing a novelty. Both regional proxies for public knowledge are positively correlated with the introduction of novelties. However, when these two regressors are considered together (column (7), only public research output (scientific publications) remains statistically significant, a finding which is consistent with Arts and Veugelers (2020). Therefore, we rely on the results in regression (4), being the most conservative estimates for the effect of co-authored publications, to infer variation in novelties and use the predicted value of this variable in the second stage (event-analysis) regression as treatment (Eq. (4)). The bottom panel of Figure 1 confirms a non-linear relationship between novelties and productivity growth, consistent with the pattern documented previously. However, in this case, the parameter size is smaller, as most of the variation in the impact of novelties is absorbed by the region-level variables included in the probit regression.

To mitigate any remaining bias potentially associated with selectivity, we replicate the Local Projection regressions by restricting the control group to firms identified as similar to the treated units using the matching procedure described above. We consider two alternative control groups, obtained using respectively ten and five nearest neighbours. The results of these event studies are illustrated in Figure 2. Notably, both the pattern and the magnitude of the estimated effects for novelties are consistent with those shown in Figure 1, falling between the corresponding parameter estimates. As expected, the

Figure 1: Event analysis (Full sample)

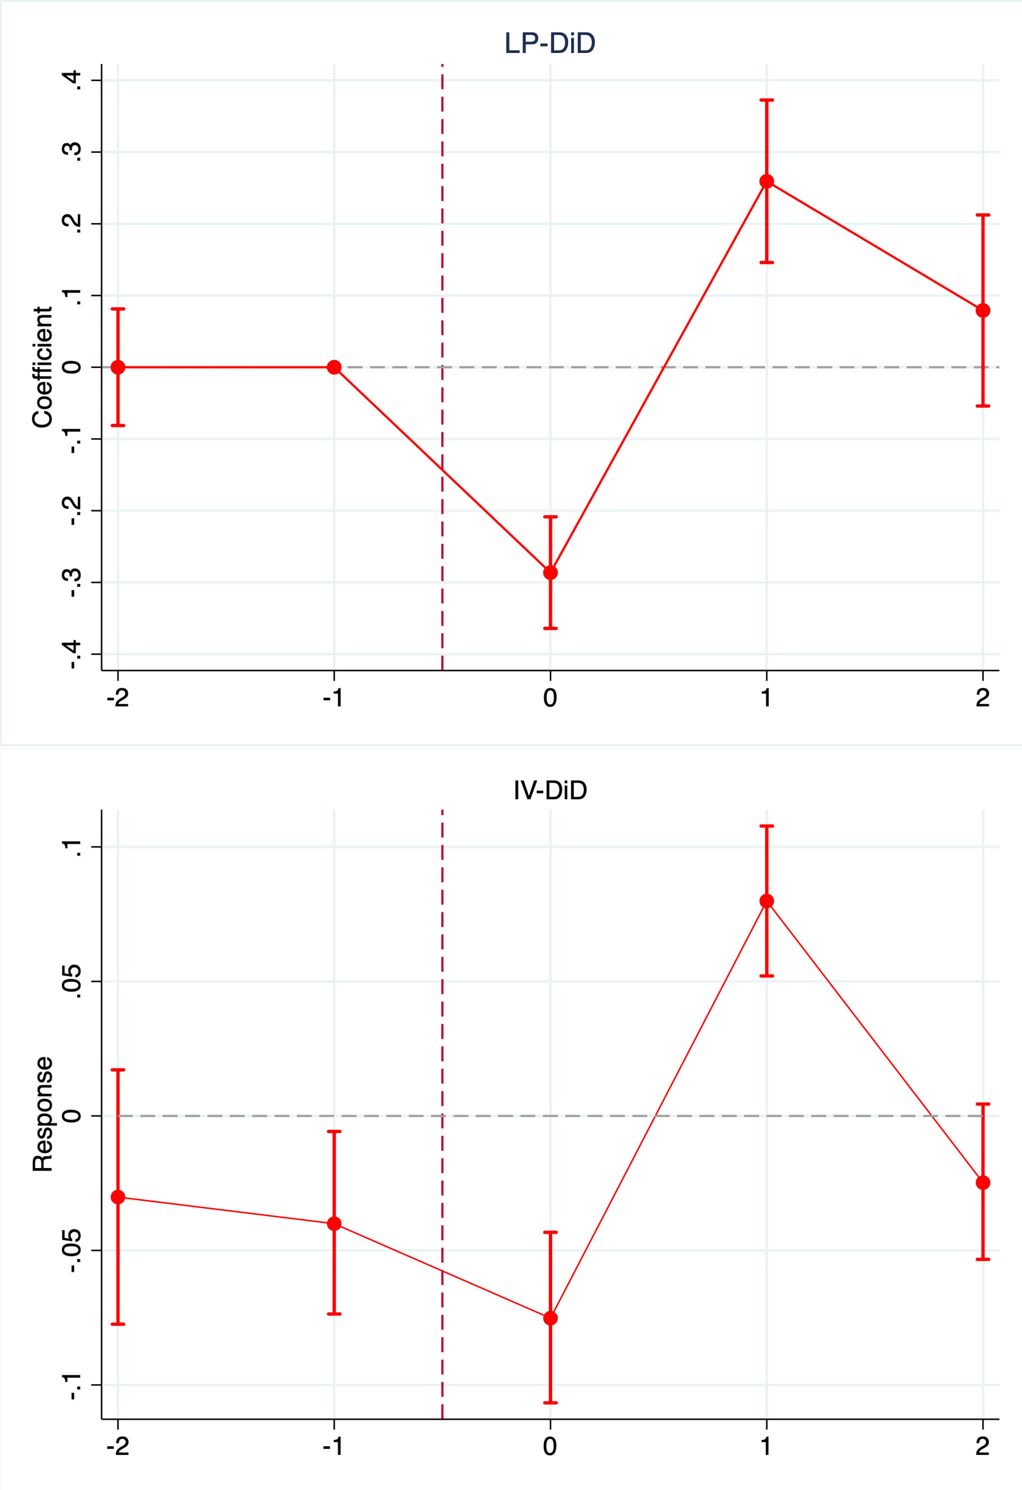

Notes: LP-DiD and IV-DiD estimates are obtained controlling for one-year lag of the dependent variable and the logs of lagged productivity levels, the capita-labor ratio and patents per worker. IV-DiD regression uses wild-bootstrapped standard errors with 1,000 replications. First-stage regression is based on pooled probit. Second-stage regression is based on a DiD specification. Bands reflect 95% confidence intervals.

precision of estimates is smaller due to a reduction in the regression sample.

Figure 2: Event analysis (Matched sample)

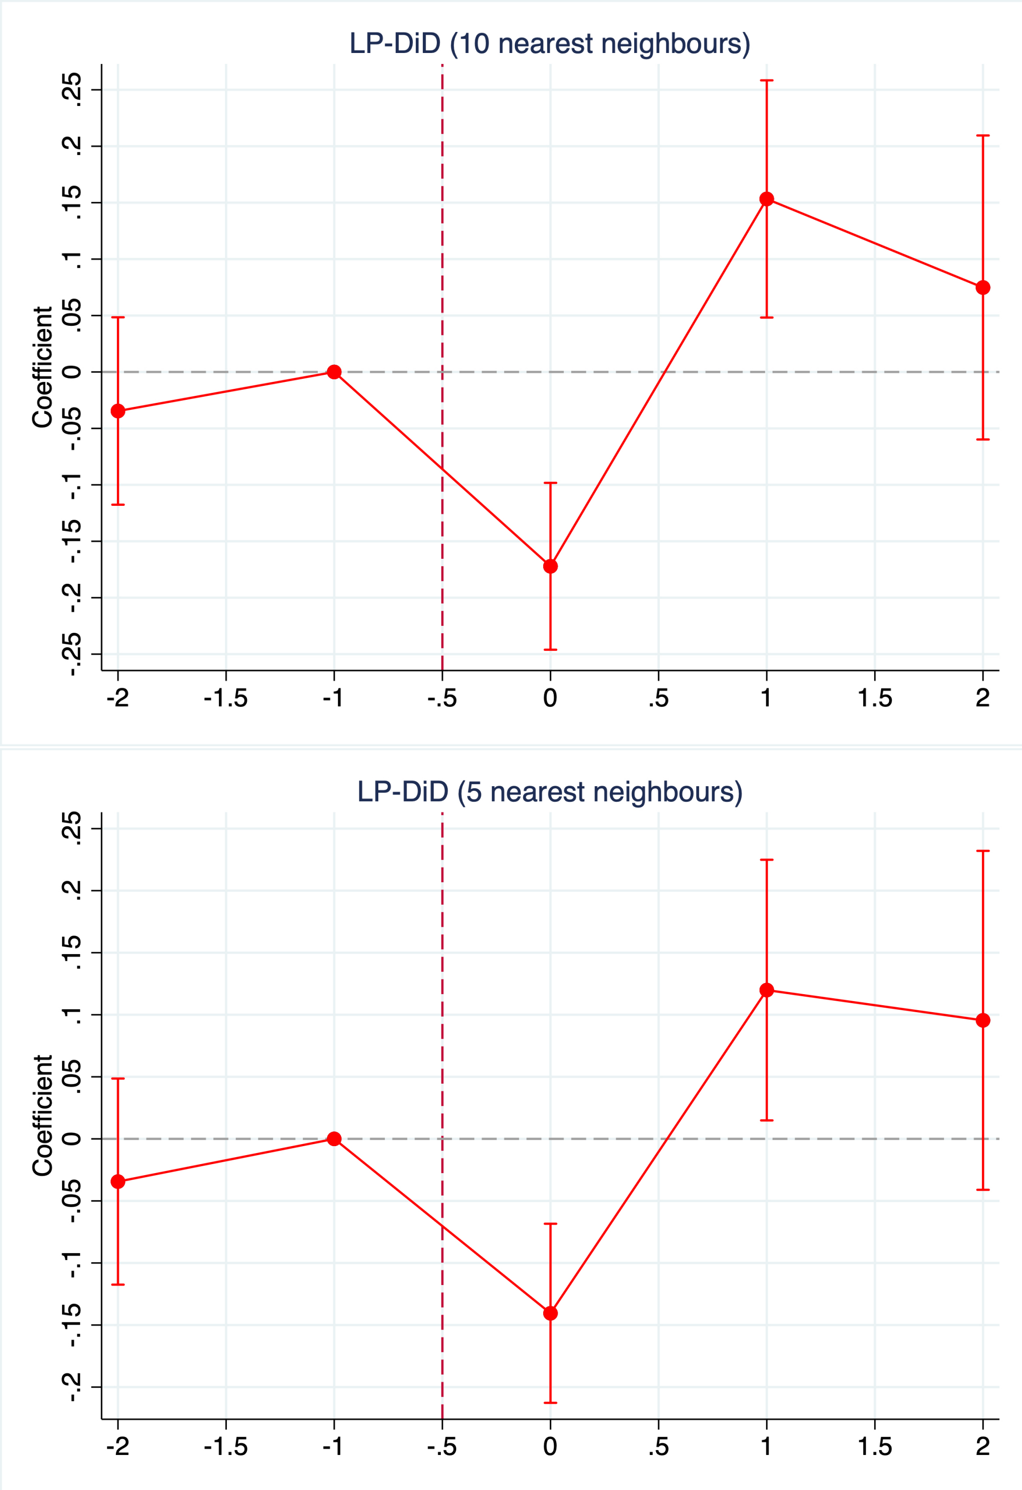

Notes: LP-DiD estimates are obtained using the samples of matched companies identified with a propensity matching procedure estimated via single-year probit regressions applying the nearest neighbour (NN) and common support. Both probit and LP-DiD regressions control for one-year lag of the dependent variable and the logs of lagged productivity levels, the capita-labor ratio and patents per worker. Bands reflect 95% confidence intervals.

Transmission channels and heterogeneity in the effects

We now turn to the key question of the mechanisms underlying both the initial penalty associated with radical innovations and the subsequent gains. To this end, we decompose the growth rate of labour productivity into the growth of value added (or sales) and growth of employment, and examine whether the dominant channel differs across sectors and firm sizes.

Table 6: **Impact of novelties on output growth and employment growth**

	(1)	(2)	(3)	(4)	(5)	(6)
	Text			**Backward cites**		
	Output growth		Employment growth	Output growth		Employment growth
	Value added	Sales		Value added	Sales	
Novelty (contemp.)	0.029	-0.005	0.436***	-0.009	-0.019	0.383***
	(0.036)	(0.025)	(0.031)	(0.034)	(0.028)	(0.033)
Novelty (cumulative)	-0.0213	0.042*	-0.093***	-0.058*	0.031	-0.092***
	(0.040)	(0.025)	(0.021)	(0.041)	(0.028)	(0.027)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	146,901	179,708	196,896	146,901	179,708	196,896
R-squared	0.509	0.511	0.722	0.509	0.511	0.721

Notes: The dependent variable is the annual rate of change of labour productivity. Parameter estimates are obtained from OLS regressions including firm and industry-by-year fixed effects. Controls (in logs) include lagged labour productivity, the capital–labour ratio, and patents per worker. Standard errors, reported in parentheses, are robust to heteroskedasticity and autocorrelation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The results on the mechanisms are reported in Table 6. They indicate that the initial negative effect is mainly driven by an increase in employment growth following the introduction of novelties, which is not matched by a proportional rise in output growth. Value-added growth appears unresponsive to the introduction of novelties when measured using patent text, while sales exhibit a modest yet statistically weak delayed reaction to radical innovation. A broadly similar pattern emerges when novelties are measured through backward citations. This pattern mirrors, in reverse, the dynamic effects found on labour productivity growth reported in Table 3. The employment-driven adjustment mechanism observed after the introduction of radical innovations suggests that firms initially expand labour inputs to reorganize production. Subsequently, as the new technology becomes fully integrated within the firm, employment reallocates across activities, expanding in new business lines while contracting in areas that are no longer competitive, with the latter effect being probably dominant. This dynamic adjustment appears to be a key characteristic of firms operating, among others, in the field of digital innovation (Gao et al., 2025).

Table 7 examines the heterogeneity in the effects of radical innovation across macro-sectors, distinguishing between firms operating in manufacturing (including mining) and in services (utilities and business services). In the short run, the negative impact of radical innovation on labour productivity growth is considerably stronger in the service sector. In the medium run, the cumulative effect of novelties is insignificant in manufacturing. While the underlying adjustment mechanisms appear broadly similar across macro-sectors, the magnitude of the employment response is markedly larger in the tertiary sector. These findings are consistent with evidence on the impact of early patenting on firm performance in Norway, where the employment response is found to exceed that of output, particularly

Table 7: **Channels by sectors**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Manufacturing				**Services**			
	Labour productivity growth	Output growth		Employment growth	Labour productivity growth	Output growth		Employment growth
		Value added	Sales			Value added	Sales	
Novelty (contemp.)	-0.232***	0.032	0.035	0.334***	-0.519***	-0.017	-0.114*	0.746***
	(0.043)	(0.037)	(0.025)	(0.031)	(0.106)	(0.090)	(0.059)	(0.066)
Novelty (cumulatvie)	0.037	-0.008	0.027	-0.068***	0.167**	-0.06	0.079	-0.148***
	(0.034)	(0.032)	(0.024)	(0.016)	(0.082)	(0.075)	(0.048)	(0.039)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Obs.	82,956	85,400	102,582	109,889	59,537	61,501	77,126	87,007
R-squared	0.588	0.530	0.506	0.708	0.586	0.496	0.515	0.736

Notes: The dependent variable is the annual rate of change of output (value added or sales) and employment. Parameter estimates are obtained from OLS regression that includes firm fixed effects and time-by-industry fixed effects. Controls (in logs) include lagged labour productivity, the capital–labour ratio, and patents per worker. Standard errors, reported in parentheses, are robust to heteroskedasticity and autocorrelation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

within services (Nilsen and Raknerud, 2024).

Table 8 shows that the effects documented above are entirely driven by small firms, with no statistically significant impact for medium-sized or large firms. For small firms, novelty primarily affects labour productivity growth through employment expansion, with a more limited contribution from output (sales) growth (cols. (3)–(4)). The short-run productivity effect is strongly negative and reflects occupational expansion. However, in the relatively longer run, radical innovation yields economically sizable gains, raising labour productivity growth by about 17 percent. These findings lend support to previous evidence pointing to a greater sensitivity of SMEs to radical innovation shocks: such firms initially incur larger adjustment costs, but also exhibit greater potential for relative long-run gains once new radical technologies are fully implemented within the organisation (Baumann and Kritikos, 2016).

7 Concluding remarks

This paper shows that radical innovations, here defined as novelties, generate significant medium-term productivity growth gains, while having a negative impact in the short run. Examining a panel of over 28,500 Italian firms observed between 2012 and 2020, we document a productivity growth trajectory consistent with the "J-curve" described by Brynjolfsson et al. (2021). After introducing a radical innovation, firms experience a short-term deceleration followed by sustained improvements, with a cumulative net effect of roughly an 8% increase in labour productivity growth. The dynamics are primarily driven by employment adjustments, which expand and then contract, a pattern most pronounced among small firms.

These results align with the idea that radical innovations are often competence-destroying (Tushman and Anderson, 1986), favouring firms capable of rapidly reconfiguring resources. Small and young firms, less constrained by legacy assets, tend to adapt more effectively to technological discontinuities (Baumann and Kritikos, 2016; Cohen and Klepper, 1996). The literature shows that the ability of these companies to convert novelties into productivity growth, however, depends on access to complementary assets such as finance, skilled labour, and collaborative R&D networks (Hall and Lerner, 2010; Grimpe and Hussinger, 2013). Public policy can strengthen these assets through targeted subsidies, cost-sharing mechanisms, and instruments that facilitate access to external knowledge, such as voucher schemes (OECD, 2010). Our findings underscore the importance of regional innovation ecosystems and the interplay between public research institutions and private firms. Managerially, firms should anticipate transitional costs and complement R&D investment with workforce training, digital capability development, and organizational redesign (Tiberius et al., 2021). Enhancing absorptive capacity through sustained collaboration with universities can accelerate post-innovation recovery (Grimpe and Hussinger, 2013). Ambidextrous strategies, balancing exploration and exploitation (O'Reilly and Tushman, 2013), can further support the transformation of novelty into long-term competitiveness, consistent with evidence on the role of intangible and human capital in enabling the diffusion of General Purpose Technologies (Brynjolfsson et al., 2021).

Methodologically, we show that NLP-based novelty measures are more informative than citation-based indicators, likely because they capture semantic distance rather than ex-post citation dynamics. Our main indicator of radical innovation, N_{NLP}, defines novelty as the maximum semantic distance from the five-year prior-art corpus, aligning with research suggesting that maximum-distance metrics may better identify conceptually exceptional inventions (Gerken and Moehrle, 2012). A key advantage is timeliness: the indicator relies solely on the information available at the priority date and can be computed immediately upon filing. Because we demonstrate that text-based novelty is predictive of firm productivity growth, this immediacy provides decision-relevant information for firms, policymakers, and investors, enabling them to form expectations about the economic significance of a newly filed technology as soon as it appears. By contrast, impact and breakthrough indicators that depend on forward citations or forward textual reuse are inherently *ex post*. They are useful for retrospective explanation, but by the time they are observable, most effects of the innovation have already materialized, limiting their informational value for time-sensitive private and public decision-making.

Despite these contributions, several limitations remain. Sectoral heterogeneity is substantial: service firms exhibit the J-curve pattern, whereas manufacturing firms show clearer short-term contractions and weaker medium-term recoveries. Our radical innovation measure relies exclusively on patent data, excluding non-patenting firms and innovations embodied in trade secrets, designs, or organisational changes. Restricting to English-language patents may also introduce selection bias, favouring internationally oriented firms. Future research could integrate indicators of non-technological innovation, including digital adoption or business model changes, using survey or firm-level data. Another limitation is our absorbing treatment specification, which classifies firms as treated from the first novelty introduction onward and ignores subsequent innovation episodes and changes in novelty intensity. This may compress multiple adjustment cycles into an average post-treatment effect, smoothing heterogeneous short-run productivity penalties and medium-run recoveries following successive novelty shocks. Moreover, the lag between patent filing and market implementation means the priority year may not align with the actual diffusion of the innovation, attenuating contemporaneous effects and shifting the observed J-curve when commercialization is delayed. Future work could address these issues with non-absorbing, stacked event-study designs that reset event time around each radical innovation episode and model treatment intensity via continuous or discretized NLP-based novelty measures or counts of novel patents, thereby capturing the full dynamics of repeated shocks of varying intensity.

Together, these findings show that integrating text mining with established economic frameworks is key to understanding how technological change drives productivity. By connecting empirical evidence, new methods, and policy insight, the study offers a comprehensive view of how radical innovation develops within firms and sectors, and how its temporal, organizational, and institutional dimensions shape the translation of breakthroughs into sustained productivity growth.

Table 8: **Channels by firm size**

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Small firms				Medium firms				Large firms			
	Labour productivity growth	Output growth		Employment growth	Labour productivity growth	Output growth		Employment growth	Labour productivity growth	Output growth		Employment growth
		Value added	Sales			Value added	Sales			Value added	Sales	
Novelty (contemp.)	-0.435***	0.059	-0.005	0.652***	-0.073*	0.001	0.008	0.068***	-0.001	0.002	-0.056	0.009
	(0.066)	(0.058)	(0.038)	(0.044)	(0.044)	(0.040)	(0.034)	(0.017)	(0.077)	(0.073)	(0.045)	(0.023)
Novelty (cumulative)	0.170***	0.040	0.091***	-0.124***	-0.040	-0.103**	-0.024	-0.022	-0.027	-0.023	-0.047	0.013
	(0.051)	(0.049)	(0.034)	(0.024)	(0.040)	(0.038)	(0.029)	(0.015)	(0.071)	(0.071)	(0.043)	(0.023)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	105,185	108,539	133,897	148,959	26,578	27,024	32,093	33,744	9,165	9,675	12,042	12,387
R-squared	0.587	0.507	0.509	0.762	0.585	0.531	0.533	0.806	0.643	0.492	0.540	0.890

Notes: The dependent variable is the annual rate of change of labour productivity, output (value added or sales) or employment. Parameter estimates are obtained from the OLS regression, including firm fixed effects and time-by-industry fixed effects. Controls (in logs) include lagged labour productivity, the capital–labour ratio, and patents per worker. Standard errors, reported in parentheses, are robust to heteroskedasticity and autocorrelation. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Authors' declaration

During the preparation of this work, the authors used ChatGPT for language editing and proofreading. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

Ahuja, G. and M. C. Lampert (2001). Entrepreneurship in the large corporation: A longitudinal study of how established firms create breakthrough inventions. *Strategic Management Journal 22*(6-7), 521–543.

Albert, M., D. Avery, F. Narin, and P. McAllister (1991). Direct validation of citation counts as indicators of industrially important patents. *Research Policy 20*(3), 251–259.

Archibugi, D., A. Filippetti, and M. Frenz (2013). Economic crisis and innovation: Is destruction prevailing over accumulation? *Research Policy 42*(2), 303–314.

Arts, S., J. Hou, and J. C. Gomez. Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measures. *50*(2), 104144.

Arts, S., J. Hou, and J. C. Gomez (2021). Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measures. *Research policy 50*(2), 104144.

Arts, S. and R. Veugelers (2020). Taste for science, academic boundary spanning, and inventive performance of scientists and engineers in industry. *Industrial and Corporate Change 29*(4), 917–933.

Ashtor, J. H. (2019). Investigating cohort similarity as an ex ante alternative to patent forward citations. *Journal of Empirical Legal Studies 16*(4), 848–880.

Athey, S. and G. W. Imbens (2022). Design-based analysis in Difference-in-Differences settings with staggered adoption. *Journal of Econometrics 226*(1), 62–79.

Audretsch, D. B. and M. P. Feldman (1996). R&D spillovers and the geography of innovation and production. *American Economic Review 86*(3), 630–640.

Balasubramanian, N. and J. Sivadasan (2011). What happens when firms patent? New evidence from U.S. economic census data. *Review of Economics and Statistics 93*(1), 126–146.

Banerjee, P. M. and B. M. Cole (2010). Globally radical technologies and locally radical technologies: The role of audiences in the construction of innovative impact in biotechnology. *IEEE Transactions on Engineering Management 58*(2), 262–274.

Baumann, J. and A. S. Kritikos (2016). The link between R&D, innovation and productivity: Are micro firms different? *Research Policy 45*(6), 1263–1274.

Bloom, N. and J. Van Reenen (2002). Patents, real options and firm performance. *Economic Journal 112*(478), C97–C116.

Bontempi, M. E. and J. Mairesse (2015). Intangible capital and productivity at the firm level: A panel data assessment. *Economics of Innovation and New Technology 24*(1-2), 22–51.

Brown, J. R., S. M. Fazzari, and B. C. Petersen (2009). Financing innovation and growth: Cash flow, external equity, and the 1990s R&D boom. *Journal of Finance 64*(1), 151–185.

Brynjolfsson, E., D. Rock, and C. Syverson (2021). The productivity j-curve: How intangibles complement General Purpose Technologies. *American Economic Journal: Macroeconomics 13*(1), 333–372.

Carpenter, M. P., F. Narin, and P. Woolf (1981). Citation rates to technologically important patents. *World Patent Information 3*(4), 160–163.

Castaldi, C., K. Frenken, and B. Los (2015). Related variety, unrelated variety and technological breakthroughs: An analysis of US state-level patenting. *Regional Studies 49*(5), 767–781.

Castellacci, F. and J. Zheng (2010). Technological regimes, Schumpeterian patterns of innovation and firm-level productivity growth. *Industrial and Corporate Change 19*(6), 1829–1865.

Chappell, N. and A. Jaffe (2018). Intangible investment and firm performance. *Review of Industrial Organization 52*(4), 509–559.

Christensen, C. M. (2013). *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston, MA: Harvard Business Review Press.

Christensen, C. M. and R. S. Rosenbloom (1995). Explaining the attacker's advantage: Technological paradigms, organizational dynamics, and the value network. *Research Policy 24*(2), 233–257.

Cohen, W. M. (2010). Fifty years of empirical studies of innovative activity and performance. In B. H. Hall and N. Rosenberg (Eds.), *Handbook of The Economics of Innovation, Vol. 1*, Volume 1 of *Handbook of the Economics of Innovation*, pp. 129–213. North-Holland.

Cohen, W. M. and S. Klepper (1996). Firm size and the nature of innovation within industries: The case of process and product R&D. *Economic Journal 106*(437), 925–951.

Crépon, B., E. Duguet, and J. Mairessec (1998). Research, innovation and productivity: An econometric analysis at the firm level. *Economics of Innovation and New Technology 7*(2), 115–158.

Criscuolo, C., D. Andrews, and P. N. Gal (2019). The best versus the rest: Divergence across firms during the global productivity slowdown. Technical report, London School of Economics and Political Science. Centre for Economic Performance, LSE.

Dahlin, K. B. and D. M. Behrens (2005). When is an invention really radical? defining and measuring technological radicalness. *Research Policy 34*(5), 717–737.

Dewar, R. D. and J. E. Dutton (1986). The adoption of radical and incremental innovations: An empirical analysis. *Management Science 32*(11), 1422–1433.

Dosi, G. (1982). Technological paradigms and technological trajectories: A suggested interpretation of the determinants and directions of technical change. *Research policy 11*(3), 147–162.

Dube, A., D. Girardi, Jordà, and A. M. Taylor (2025). A Local Projections approach to Difference-in-Differences. *Journal of Applied Econometrics https://doi.org/10.1002/jae.70000*.

Fernald, J., R. Inklaar, and D. Ruzic (2025). The productivity slowdown in advanced economies: Common shocks or common trends? *Review of Income and Wealth 71*(1), e12690.

Fleming, L. (2001, January). Recombinant uncertainty in technological search. *Management Science 47*(1), 117–132.

Fronzetti Colladon, A., B. Guardabascio, and F. Venturini (2025). A new mapping of technological interdependence. *Research Policy 54*(1), 105126.

Funk, R. J. and J. Owen-Smith (2017). A dynamic network measure of technological change. *Management Science 63*(3), 791–817.

Gao, J., Z. Li, T. Nguyen, and W. Zhang (2025). Digital transformation and enterprise employment. *International Review of Economics & Finance 99*, 104036.

Gerken, J. M. and M. G. Moehrle (2012, June). A new instrument for technology monitoring: Novelty in patents measured by semantic patent analysis. *Scientometrics 91*(3), 645–670.

Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature 28*(4), 1661–1707.

Grimpe, C. and K. Hussinger (2013). Formal and informal knowledge and technology transfer from academia to industry: Complementarity effects and innovation performance. *Industry and Innovation 20*(8), 683–700.

Gómez, J., I. Salazar, and P. Vargas (2020). The role of extramural R&D and scientific knowledge in creating high novelty innovations: An examination of manufacturing and service firms in Spain. *Research Policy 49*(8), 104030.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2005). Market value and patent citations. *RAND Journal of Economics 36*(1), 16–38.

Hall, B. H. and J. Lerner (2010). The financing of R&D and innovation. In *Handbook of the Economics of Innovation*, Volume 1, pp. 609–639. Elsevier.

Hall, B. H., J. Mairesse, and P. Mohnen (2010). Measuring the returns to R&D. In *Handbook of the Economics of Innovation*, Volume 2, pp. 1033–1082. Elsevier.

Hall, B. H. and M. Trajtenberg (2004). Uncovering GPTs with patent data. NBER Working Paper 10901, National Bureau of Economic Research. Available at SSRN: https://ssrn.com/abstract=618583.

Henderson, R. (1993). Underinvestment and incompetence as responses to radical innovation: Evidence from the photolithographic alignment equipment industry. *RAND Journal of Economics 24*(2), 248 – 270.

Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics 108*(3), 577–598.

Kauffman, S. A. (2000). *Investigations*. Oxford University Press.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021). Measuring technological innovation over the long run. *American Economic Review: Insights 3*(3), 303–320.

Lanjouw, J. O. and M. Schankerman (2001). Characteristics of patent litigation: A window on competition. *RAND Journal of Economics*, 129–151.

Lanjouw, J. O. and M. Schankerman (2004). Patent quality and research productivity: Measuring innovation with multiple indicators. *Economic Journal 114*(495), 441–465.

Lipsey, R. G., K. I. Carlaw, and C. T. Bekar (2005, November). Technology and technological change. In *Economic Transformations: General Purpose Technologies and Long-Term Economic Growth*, Chapter 4, pp. 85–130. Oxford University Press.

Marioni, L. d. S., A. Rincon-Aznar, and F. Venturini (2024). Productivity performance, distance to frontier and AI innovation: Firm-level evidence from Europe. *Journal of Economic Behavior Organization 228*, 106762.

Mokyr, J. (1990). Punctuated equilibria and technological progress. *American Economic Review 80*(2), 350–354.

Nasirov, S. and C. Castaldi (2025). Servitization as a strategy for diffusing radical technologies: An analysis of U.S. top corporate R&D investors. *Research Policy 54*(1), 105123.

Nelson, R. R. and S. G. Winter (1982). *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard University Press.

Nilsen, O. A. and A. Raknerud (2024). Dynamics of first-time patenting firms. *Research Policy 53*(8), 105054.

OECD (2010). *Innovation Vouchers*. Paris: Organisation for Economic Co-operation and Development.

O'Reilly, C. A. and M. L. Tushman (2013). Organizational ambidexterity: Past, present, and future. *Academy of Management Perspectives 27*(4), 324–338.

Petralia, S. (2020). Mapping General Purpose Technologies with patent data. *Research Policy 49*(7), 104013.

Rice, M., D. Kelley, L. Peters, and G. Colarelli O'Connor (2001). Radical innovation: Triggering initiation of opportunity recognition and evaluation. *R&D Management 31*(4), 409–420.

Schoenmakers, W. and G. Duysters (2010). The technological origins of radical inventions. *Research Policy 39*(8), 1051–1059.

Schumpeter, J. A. (1934). *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest, and the Business Cycle*. Cambridge, MA: Harvard University Press.

Shane, S. (2001). Technological opportunities and new firm creation. *Management Science 47*(2), 205–220.

Squicciarini, M., H. Dernis, and C. Criscuolo (2013). Measuring patent quality: Indicators of technological and economic value. Technical Report 2013/03, OECD Publishing, Paris.

Tiberius, V., H. Schwarzer, and S. Roig-Dobón (2021). Radical innovations: Between established knowledge and future research opportunities. *Journal of Innovation and Knowledge 6*(4), 250–262.

Trajtenberg, M. (1990). A penny for your quotes: Patent citations and the value of innovations. *RAND Journal of Economics 21*(1), 172–187.

Tria, F., V. Loreto, V. D. P. Servedio, and S. H. Strogatz (2014). The dynamics of correlated novelties. *Scientific reports 4*(1), 5890.

Tushman, M. L. and P. Anderson (1986). Technological discontinuities and organizational environments. *Administrative Science Quarterly 31*(3), 439–465.

Ugur, M., E. Trushin, E. Solomon, and F. Guidi (2016). R&D and productivity in OECD firms and industries: A hierarchical meta-regression analysis. *Research Policy 45*(10), 2069–2086.

Usher, A. P. (1954). *A History of Mechanical Invention*. Cambridge, MA: Harvard University Press.

Verhoeven, D., J. Bakker, and R. Veugelers (2016). Measuring technological novelty with patent-based indicators. *Research Policy 45*(3), 707–723.

Verluise, C., G. Cristelli, K. Higham, and G. de Rassenfosse (2025). Beyond the front page: In-text citations to patents as traces of inventor knowledge. *Strategic Management Journal*.

Online Appendix

Table A.1: Descriptive Statistics: Backward citations-based novelty

	Treated (2,377 firms)			Control (26,234 firms)		
	Mean	Median	SD	Mean	Median	SD
Labour productivity growth (%)	-0.17	-0.32	75.45	-0.04	-0.05	88.96
Controls						
Capital per worker (mln)	28.22	8.58	227.54	66.92	6.30	1435.26
Patents per worker	0.17	0.03	2.59	0.41	0.06	14.88
# of innovators by region-sector	9.06	0	26.06	0.74	0	7.79
# of innovators by region-sector	224.22	66	333.77	139.47	26	258.80
Cash-flow ratio	8.07	0.073	539.41	24.12	0.05	4043.67
Intangible Ratio	0.22	0.09	0.28	0.18	0.05	0.26
Debt Ratio	0.70	0.73	0.20	0.71	0.75	0.23
Employment growth (%)	0.11	0	39.79	0.15	0	46.65
Value Added growth (%)	-0.05	-0.14	62.37	0.11	-0.14	74.67

Table A.2: Baseline estimates Novelty Mean

	(1)	(2)	(3)	(4)
Novelty (post × treatment)	-0.057***	-0.076***	-0.080***	-0.060***
	(0.023)	(0.026)	(0.026)	(0.025)
Lagged productivity (log)	-0.010***	-0.010***	-0.010***	-0.010***
	(0.000)	(0.000)	(0.000)	(0.000)
Capital per worker (log)	0.132***	0.129***	0.129***	0.165***
	(0.005)	(0.004)	(0.004)	(0.006)
Patents per worker (log)	0.018***	0.020***	0.020***	0.021***
	(0.001)	(0.001)	(0.001)	(0.001)
# of innovators by region-sector (log)			0.000***	
			(0.000)	
# of patents by region-sector (log)			-0.000	
			(0.000)	
Cash flow ratio				-0.081***
				(0.001)
Intangible ratio				-0.008***
				(0.002)
Debt ratio				0.165***
				(0.012)
Labor productivity measure	Sales p.w.	Value added p.w.	Value added p.w.	Value added p.w.
Observations	238,873	189,802	189,802	147,338
R-squared	0.564	0.560	0.560	0.610

Notes: The dependent variable is the annual rate of change of labour productivity. Parameter estimates are obtained from an OLS regression including firm and industry-by-year fixed effects. For each firm, the novelty indicator is computed as the mean novelty of its patents relative to technologies patented in the preceding five years. Standard errors are robust to heteroskedasticity and autocorrelation in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A.3: Dynamic effects Novelty Mean

	(1)	(2)	(3)	(4)	(5)
Novelty (post $\times$ treatment)	-0.076***	-0.421***	-0.419***	-0.418***	-0.515***
	(0.026)	(0.041)	(0.049)	(0.056)	(0.056)
Novelty $t-1$		0.450***	0.416***	0.445***	0.605***
		(0.039)	(0.052)	(0.056)	(0.059)
Novelty $t-2$			0.068**	-0.036	0.017
			(0.031)	(0.040)	(0.038)
Novelty $t-3$				0.104***	0.084***
				(0.034)	(0.031)
Novelty (cumulative)		0.028	0.064*	0.095**	0.191***
		(0.026)	(0.034)	(0.044)	(0.039)
Controls	Yes	Yes	Yes	Yes	Yes
Labor productivity measure	Value added p.w.	Value added p.w.	Value added p.w.	Value added p.w.	Sales p.w.
Observations	189,802	189,802	166,431	142,493	179,304
R-squared	0.560	0.560	0.571	0.585	0.586

Notes: The dependent variable is the annual rate of change of labour productivity. Parameter estimates are obtained from an OLS regression including firm and industry-by-year fixed effects. For each firm, the novelty indicator is computed as the mean novelty of its patents relative to technologies patented in the preceding five years.Controls (in logs): lagged labor productivity, capital-labor ratio, patents per worker. Standard errors are robust to heteroskedasticity and autocorrelation in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A.1: Event analysis using mean novelty (Full sample)

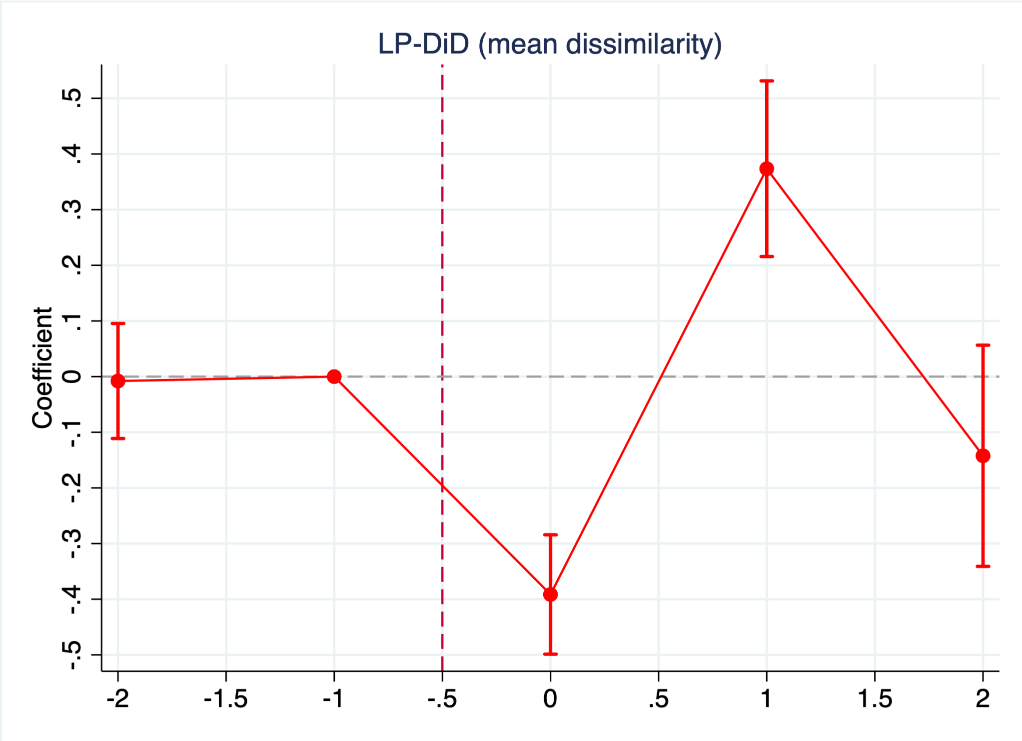

Notes: LP-DiD estimates are obtained controlling for one-year lag of the dependent variable and the logs of lagged productivity levels, the capita-labor ratio and patents per worker (all in logs). For each firm, the novelty indicator is computed as the mean novelty of its patents relative to technologies patented in the preceding five years. Bands reflect 95% confidence intervals.