

Big firms, small firms, and innovation: the German coal and steel industry, 1928–1941

Julian Lohf*

20th May 2026

Abstract

Did large firms drive technological progress, as Joseph Schumpeter famously claimed? This article revisits the Schumpeterian hypothesis using a newly constructed dataset of more than ten thousand patent applications published in the trade journal *Stahl und Eisen* between 1928 and 1943, linked to balance-sheet information for some six hundred firms in the German coal and steel industry. The setting is unusual: the same firms can be observed as the institutional environment shifts from the market economy of the late Weimar Republic to the autarky and rearmament regime of National Socialism. The dependent variable is a count of patent applications, and the model is a fixed-effects Poisson estimated on a panel that includes the firm-years without any application. The relationship between firm size and innovation was not monotonic: the smallest firms and the largest concerns applied for patents most intensively, while medium-sized firms lagged behind. The pattern is sharp in the years before the Four-Year Plan of 1936 and cannot be detected after it, when the spread of secret patenting makes the published record an unreliable guide. A correctly constructed index of regional patent concentration enters in the Schumpeterian direction but is significant only in some specifications, so the evidence on market structure remains suggestive rather than settled. The evidence as a whole qualifies rather than confirms the Schumpeterian hypothesis.

Keywords: innovation; firm size; patents; Schumpeterian hypothesis; coal and steel industry; Germany; National Socialism; cliometrics

JEL classification: N64; O31; L11

*[institutional affiliation]. Working draft — comments welcome.

I Introduction

Few propositions in the economics of innovation have proved as durable, or as contested, as Joseph Schumpeter's contention that large firms are the principal engines of technological progress. In *Capitalism, Socialism and Democracy*, Schumpeter argued that the most conspicuous improvements in economic life could be traced not to small firms operating under competitive conditions but to the large concerns that dominated their markets (Schumpeter, 1942). The claim — that firm size and market power foster innovation — became known as the Schumpeterian hypothesis, and for more than seven decades it has framed debate over the relationship between industrial structure and technical change (Kamien and Schwartz, 1975; Gilbert, 2006). Yet despite a large empirical literature, the question Schumpeter posed remains unresolved.

The empirical record is genuinely mixed. Studies have variously reported that innovation rises with firm size, that it falls, and that it bears no systematic relationship to size at all (Mansfield, 1963; Acs and Audretsch, 1987, 1988). One reason for this lack of consensus is that results are unusually sensitive to the data on which they rest. The measure of innovation, the length of the observation window, the geographical scope of the sample, and the treatment of small and non-incorporated firms all shape the conclusions reached. As Kamien and Schwartz (1975) noted in an early survey, the selection of inputs, outputs and periods has been driven largely by what data happen to be available, so that findings are difficult to compare across studies and settings.

This article reconsiders the Schumpeterian hypothesis using a new historical dataset for a single industry observed in depth: the German coal and steel industry — the *Montanindustrie* — between 1928 and 1941. The core of the dataset comprises more than ten thousand patent applications published in the leading trade journal *Stahl und Eisen*, hand-collected for the years 1928 to 1943 and linked, firm by firm, to balance-sheet information for some six hundred enterprises drawn from the *Handbücher der deutschen Aktiengesellschaften* and supplementary company registers. Equity capital serves as the measure of firm size; the firm-level data are combined with regional information on human capital and urbanisation. Three features distinguish the dataset. First, it records patent *applications* rather than grants: an application marks the point at which a firm commits to protecting the output of its research, and so tracks inventive effort more directly than the later, administratively filtered grant (Grupp, 1997). Second, because the entire observation window lies within a single technological paradigm and a single, unchanged system of patent classification — the *Deutsche Patentklassifikation* — inventive activity is measured on a consistent basis throughout. Third, the dataset reaches the very smallest firms and those organised as private limited companies, partnerships and other non-corporate forms, which the standard joint-stock-company sources systematically omit.

The period is more than a convenient sample. Between 1928 and 1941 the same firms

can be observed as the institutional environment around them was transformed. The late Weimar Republic’s comparatively liberal market economy gave way, after the Depression, to the controlled foreign trade, autarky and accelerating rearmament of the National Socialist regime, institutionalised in the New Plan of 1934, the Four-Year Plan of 1936 and the creation of the state-owned *Reichswerke* in 1937. The *Montanindustrie* thus offers an unusual historical laboratory: it allows the size–innovation relationship to be examined not as a fixed structural fact but as something that may itself respond to the economic order in which firms operate. Did the relationship change as coordination shifted from the market towards command and rearmament?

The central findings can be stated briefly. The relationship between firm size and innovation was not monotonic. Both the smallest firms and the largest concerns applied for patents most intensively, while medium-sized firms were the least inventive — a pattern at odds with the simple Schumpeterian claim that innovation rises with size, but equally at odds with the contrary view that small, competitive firms alone drive technical progress. The size gradient did not steepen with rearmament; if anything it flattened after 1936, a pattern that the spread of secret patenting during the war makes hard to interpret. The market-structure variables prove fragile: two measures of regional concentration enter with opposite signs, so that no firm conclusion about competition and innovation can be drawn, and the apparent association with regional human capital does not survive a proper calculation of the standard errors. Finally, the size–innovation relationship differed across the sectors of the industry.

In revisiting the Schumpeterian hypothesis in this setting, the article speaks to three bodies of work: the long-running debate on firm size and innovation; the economic history of innovation in Germany, which has so far concentrated on the *Kaiserreich* and the Weimar years (Streb et al., 2006; Degner, 2012); and the business and economic history of the German coal and steel industry under National Socialism. The remainder of the article is organised as follows. Section II sketches the historical setting. Section III situates the study within the empirical literature on firm size and innovation. Section IV describes the data and Section V the empirical strategy. Section VI presents the results; Section VII discusses robustness and the limits of the evidence; and Section VIII concludes.

II The German coal and steel industry, 1928–1941

The German coal and steel industry that entered the period under study was already highly concentrated. The years after the First World War combined rapid technical change with severe political and economic disruption (Borchardt, 1985). The Treaty of Versailles cost Germany economically important territories, its merchant fleet and a substantial body of patents, while reparations and, from 1923, hyperinflation undermined financial stability. Paradoxically, inflation also encouraged consolidation in heavy industry: the

Reich compensated firms for plant lost through territorial cessions, devalued wages acted as an implicit subsidy, and currency depreciation made German exports cheap. The most prominent product of this consolidation was the Vereinigte Stahlwerke AG, formed in 1926 as successor to the Stinnes group and, with that, the largest steel concern in Europe and the second largest in the world (Fiereder, 1983). By the late 1920s the six largest concerns — Vereinigte Stahlwerke, Krupp, Gutehoffnungshütte, Hoesch, Klöckner and Mannesmann — accounted for roughly four-fifths of German pig-iron output, and cartels and syndicates governed much of the rest.

The crisis that began in 1929 brought this order to an end. In Germany it struck a business cycle that had already passed its peak in 1928; the collapse of the Vienna Creditanstalt in 1931 added a banking crisis to an economy structurally weakened by the earlier destruction of capital. Chancellor Brüning's deflationary policy failed to arrest the downturn. Internationally, the Depression dismantled what remained of the liberal trading order: after the United States enacted the Smoot–Hawley tariff in 1930, a wave of retaliatory protection followed, and world trade fell by more than three-fifths between 1929 and 1932 (Capie, 2013). Germany increasingly responded by withdrawing from the world market and building bilateral trading relationships of its own.

The National Socialist seizure of power in 1933 brought not only a political but an economic transformation. Hjalmar Schacht, as economics minister and Reichsbank president, designed the institutional framework. His New Plan of 1934 subordinated foreign trade to administrative control: every request for foreign exchange was assessed for its economic necessity, and imports were steered towards countries with which the Reich held clearing agreements (Kopper, 2006). The policy conserved scarce currency at the cost of dearer and less flexible trade, and pushed the economy towards autarky. Rearmament was financed off the books through the so-called MeFo bills, a device that kept the true scale of military spending out of the Reich's published accounts. The regime also reshaped company law: the reform of 1934 set a minimum capital of RM 500,000 for joint-stock companies and compelled small joint-stock firms to convert or dissolve, while the tax system was tilted in favour of partnerships and other personal-liability forms. The effect was substantial — the number of joint-stock companies fell from around 8,000 in 1935 to 5,350 by the end of 1939. For a study of firm size and innovation this institutional detail matters, since it means that the population of firms cannot be observed through the joint-stock form alone.

The Four-Year Plan of 1936 made the priority explicit: the German economy, in Hitler's words, was to be fit for war within four years. The binding constraint was raw materials, above all iron ore. German imports of iron ore tripled between 1931 and 1938 while exports all but ceased; domestic ores were abundant but low-grade, with an iron content far below that of the imported Swedish ore on which the Ruhr concerns preferred to rely (Donges, 2014). When the private industry proved reluctant to invest in

smelting low-grade domestic ore, the regime acted directly. In 1937 it founded the state-owned Reichswerke “Hermann Göring”, concentrated on the new ore fields of Salzgitter, and compelled the private Ruhr concerns to surrender their ore holdings to it; the new concern’s nominal capital was raised from RM 5 million to RM 400 million within a year (Mollin, 1988). By 1939 the armed forces accounted for more than half of domestic steel consumption.

Beneath these aggregate shifts the industry remained structurally diverse. Its characteristic unit was the vertically integrated *Verbund* concern, which combined coal mining, coke production, iron and steel making, downstream processing and even power and gas supply within a single organisation. Yet the large concerns had reached this common form by markedly different routes — the Gutehoffnungshütte from iron smelting, Krupp from steel casting, Mannesmann from tube production, Klöckner partly from trade — and they differed correspondingly in the balance of mining, steelmaking and processing activities they contained. Around and below them stood a wide range of downstream and supplying sectors: machine-building, metal-working, the chemical industry and electricity generation. This heterogeneity matters for what follows, because it implies that the relationship between firm size and innovation need not be uniform across the industry.

Taken together, the period spans a pronounced change of economic regime: from the market-coordinated — if already heavily cartelised — industry of the late Weimar years, through the Depression, to the autarkic and rearmament-driven command economy of the later 1930s. The empirical analysis exploits this by separating the years before and after the Four-Year Plan of 1936, and by allowing the size–innovation relationship to differ across the sectors just described.

III Firm size and innovation: the empirical debate

The intellectual starting point is Joseph Schumpeter. In his *Theorie der wirtschaftlichen Entwicklung* of 1911 he had placed the dynamic, discontinuous appearance of innovations, carried by entrepreneurs, at the centre of economic development; in *Capitalism, Socialism and Democracy* of 1942 he drew the more provocative conclusion that the large, market-dominating concern, rather than the small competitive firm, was the characteristic site of technical progress (Schumpeter, 1911, 1942). Two propositions are conventionally distilled from his argument: that monopoly power fosters innovation, and that large firms are more innovative than small ones (Frisch, 1993). This article tests the second. The empirical literature addressing it is usually divided into input-oriented studies, which measure research effort through R&D expenditure or personnel, and output-oriented studies, which count innovative results, most often patents. The present study belongs to the second group.

That literature has not converged. Among the early output-oriented studies, Phillips

(1956) examined twenty-eight manufacturing industries between 1899 and 1939 and found no support for the view that industries populated by many small firms were technologically more advanced. Mansfield (1963), studying the largest firms in the American steel, petroleum and coal industries, reached a result of particular relevance here: in iron and steel — unlike the other two industries — it was the smaller firms that proved the more innovative. Studies of the pharmaceutical industry returned mixed verdicts: Comanor (1965) found cost disadvantages rising with size alongside a positive direct size effect, Comanor and Scherer (1969) a strongly significant link between size and patent counts, and Vernon and Gusen (1974) only a direct positive effect of size. Other studies pointed away from any simple positive relationship. Smith (1974) found that the very largest American electricity producers were less innovative than merely large ones; Neumann et al. (1982) reported a significant negative size effect in West German industry; and Pavitt et al. (1987), examining innovations in the United Kingdom between 1945 and 1983, found innovative intensity highest at both ends of the size distribution — among the very large firms and among the small and medium-sized. Acs and Audretsch (1987, 1988) concluded that the question is not which firm size innovates more, but under which conditions large or small firms enjoy a relative advantage — a reading consistent with the idea that the size–innovation relationship is regime-dependent. Further studies confirmed the diversity of findings: a positive size effect in some settings (Sveikauskas and Sveikauskas, 1982; Pagano and Schivardi, 2003), an under-proportionate one in high-technology industries (Audretsch and Acs, 1991), and size effects contingent on firms’ networks (Rogers, 2004). Later work continued to emphasise competition and concentration over sheer size (Smythe, 2001; Nicholas, 2003), and stressed that the relationship varies across industries and between product and process innovation (Vaona and Pianta, 2008).

A second, smaller body of work has examined innovation in Germany historically. Streb et al. (2006) showed that technical progress in the *Kaiserreich* advanced in waves, and that inter-industry knowledge transfer between technologically and geographically related industries was an important source of inventive activity, with distinct technology regions emerging across the Reich. Baten et al. (2007) found a significant positive effect of firm size on innovativeness among Baden firms around 1900. The closest predecessor to the present study is Degner (2012), who examined more than 1,400 German firms between 1877 and 1932 using the Baten–Streb patent database. Degner counted as innovations only those patents whose holders renewed protection for at least ten years, reasoning that sustained renewal distinguishes a commercially successful invention from a mere invention; on that demanding measure he found that essentially none of the variables he examined had a discernible effect.

The present study departs from this predecessor in ways designed to meet a long-standing methodological criticism. Kamien and Schwartz (1975), surveying the field, observed that the choice of innovation measures and observation periods had been driven

largely by data availability, lending the literature a degree of arbitrariness. This article responds on three fronts. First, it measures innovation by patent *applications* rather than by long-renewed grants. Because the cost of renewing a patent rises with its age, a measure based on long-held patents systematically under-represents small firms, for which renewal is the riskier commitment; counting applications avoids this bias and stays closer to the firm-level decision to innovate. Second, the entire observation window lies within a single technological paradigm and a single system of patent classification, so that inventive activity is measured consistently over time and the “arbitrary window” objection loses much of its force. Third, by concentrating on one industry — in which, moreover, innovation was overwhelmingly incremental — the study can take seriously the heterogeneity of firms and sectors that cross-industry studies must average away. The question it puts to this evidence is therefore a focused one: in the German coal and steel industry between 1928 and 1941, were large firms more innovative than small ones, and did the answer change as the economic regime changed?

IV Data

The empirical analysis rests on a dataset assembled specifically for this study. It combines three layers of evidence: the patent applications that measure inventive activity, the firm characteristics that measure size, and the regional variables that describe the innovation system within which firms operated. Each layer is described in turn.

Inventive activity is measured by published patent applications. The source is the trade journal *Stahl und Eisen — Zeitschrift für das deutsche Eisenhüttenwesen*, which has appeared since 1881 and which reproduced, week by week, all patent applications, utility models and granted patents issued by the Reich Patent Office that concerned the coal and steel industry. From roughly eight hundred weekly issues between 1 January 1928 and 20 May 1943, the dataset records more than ten thousand published patent applications, hand-collected from the digitised holdings of the Silesian University of Technology. Each application carries the names and locations of the inventors and of the applicant firm, a short description of the invention, and a patent-class code. The decision to count applications rather than granted patents follows Grupp’s argument that an application marks the point at which a firm commits to protecting the output of its own research, whereas the grant adds only legal, not economic, information; from the standpoint of the inventing firm, the application is the better proxy for inventive effort (Grupp, 1997).

One point about timing must be made explicit, since the analysis turns on it. What the Patent Office announced week by week, and what *Stahl und Eisen* reproduced, was not the filing of an application but its *Bekanntmachung*: the official announcement, after a preliminary examination, at which the application was laid open for public inspection. Under the patent law in force between the wars, that announcement followed filing only

once the examination had run its course — an interval that varied from case to case but averaged roughly eighteen months. This is not the fixed, statutory eighteen-month publication term of the present-day German system, which dates only from the reform of 1968; in the inter-war years the lag was simply processing time. Because the firm and regional characteristics that bear on an inventive decision are best measured when that decision was taken — at filing — rather than a year and a half later, each observation is dated back by this average interval, and all explanatory variables are measured at that earlier point. The adjustment is an approximation: the true lag differed across applications, so the reconstructed year of filing carries some measurement error, which adds to the limitations discussed in Section VII.

Patents are, of course, an imperfect measure of innovation. Not every invention is patented — Mansfield put the propensity to patent in American industry at between 50 and 86 per cent — and patents differ enormously in economic value (Mansfield, 1986; Pakes and Griliches, 1980). Yet, as Griliches concluded in his survey of the field, patent statistics remain a uniquely detailed and accessible record of technical change, with no comparable alternative (Griliches, 1990). Two features of the present data limit the usual concerns. Confining the count to one trade journal and to a single, stable classification system — the *Deutsche Patentklassifikation*, introduced in 1877 and used essentially unchanged between 1928 and 1968 — keeps the measure consistent and genuinely industry-specific. A second feature points the other way and must be stated plainly. From the first day of the war, a decree of 1 September 1939 allowed inventions important to national defence to be patented in secret, without publication; as the war continued the Reich increasingly took over such applications as trustee, and from 1943 the armaments minister wound the patent office down, the last announcement appearing on 20 May 1943. The dataset therefore probably understates inventive activity in the later war years, a point returned to in Section VII.

Inventive activity was highly concentrated among a small number of firms. Table 1 lists the most prolific applicants: a single firm, the machine-builder Demag, accounts for well over five per cent of all published applications, and the nine firms shown together account for more than a third. This “innovative few” pattern is also visible in the distribution of the outcome variable. Figure 1 plots the density of the (log) number of applications per firm and year: the distribution is strongly right-skewed, with a large mass of firm-years close to zero and a long thin tail of highly inventive observations — a shape that, as Section V explains, shapes the choice of estimator.

Because each application records the inventor’s place of residence and the firm’s location, inventive activity can be mapped across space. The regional units are the administrative districts of the Reich (*Regierungsbezirke*), the Austrian *Reichsgaue* and the *Oberlandratsbezirke* of the Protectorate of Bohemia and Moravia, which together correspond roughly to today’s NUTS-2 level. Figure 2 maps applications by firm location

Table 1: The “innovative few”: firms with the most published patent applications, 1928–1943.

Firm	Applications	Share (%)	Cumulative (%)
Demag AG	358	5.65	5.65
Siemens-Schuckertwerke AG	276	4.35	10.00
Fried. Krupp AG	275	4.34	14.34
Vereinigte Stahlwerke AG	272	4.29	18.63
Dr. C. Otto & Comp. GmbH	269	4.24	22.87
Fried. Krupp Grusonwerk AG	269	4.24	27.12
Schloemann AG	223	3.52	30.64
Allgemeine Elektrizitäts-Gesellschaft	189	2.98	33.62
Mannesmannröhrenwerke AG	183	2.89	36.50
<i>(remaining firms account for the balance.)</i>			




Figure 1: Density of the log number of published patent applications per firm-year.

and by inventor residence across the Reich in its borders of 1937. Both maps show the same pronounced clustering: inventive activity was concentrated overwhelmingly in the Ruhr district, with a secondary concentration around Berlin and scattered smaller centres elsewhere.

Firm size is measured by equity capital. The choice of a capital-based measure follows a long line of argument — from Sombart to Busse von Colbe — that capital is the most reliable single indicator of the size of an enterprise, more so than employment, which reveals nothing about the capital a firm commands (Sombart, 1927; Busse von Colbe, 1964). Equity capital is preferred to nominal share capital because the coal and steel industry contained undercapitalised companies whose disclosed reserves far exceeded their nominal capital, so that share capital alone would misstate their true size. The firm data are drawn principally from the *Handbücher der deutschen Aktiengesellschaften* for 1933–1944 — some sixty volumes, each covering between two and five thousand firms




Figure 2: Published patent applications by firm location (left) and by inventor residence (right), at district level, German Reich in the borders of 1937.

— together with the companion *Handbuch der großen GmbH, KG, OHG, Einzelfirmen*, which records firms of other legal forms with a capital of at least RM 500,000, and supplemented by company reports and the press archive of the Hamburg Institute for the World Economy. The resulting dataset covers roughly six hundred firms that filed patent applications relevant to the coal and steel industry; the smallest recorded equity capital is RM 5,000, an indication of how far down the size distribution the sources reach. Firms are sorted into six size classes, from micro-firms (class 1, capital below RM 500,000) to the great concerns (class 6, capital above RM 100 million), and the size variable enters the analysis as the natural logarithm of the size class. Firm age is measured treating mergers as new foundations; the oldest firm in the data, an ironworks tracing its origins to the fifteenth century, is over 450 years old.

One feature of the firm identification should be made explicit, because it bears on what “size” means here. Members of the same group are recorded as separate firms wherever they applied for patents under their own name and kept their own balance sheet, so that, for instance, Fried. Krupp AG and the Krupp Grusonwerk appear as distinct units. This follows the legal and accounting reality of the period — the subsidiary was a company in its own right — but it means that the size variable is the size of the legally distinct enterprise, not of the concern behind it. A concern-level aggregation, which would reclassify a number of mid-sized subsidiaries into the largest class, is left to future work; the analysis below should be read as a statement about firms as legally constituted, not about concerns.

Figure 3 shows the combined equity capital of the sample firms over the period: it falls into a trough in the mid-1930s and then rises steeply as rearmament gathered pace.

The third layer of the dataset describes the regional innovation system. Regional variables are built from the Reich statistical yearbooks, from Titze’s data handbook of German educational history and the associated HISTAT database, and from the *Handbücher* (Titze, 1987; Grupp et al., 2002). Market structure is captured by a single index, motivated by the inverted-U argument of Aghion et al. (2005): *HHI*, the concentration of coal-and-steel patent applications across the firms of a region, which measures how far




Figure 3: Combined equity capital of the sample firms, 1928–1941.

a region’s inventive activity rested on a few dominant firms. For a region i it takes the familiar Herfindahl–Hirschman form

$$HHI_i = \sum_{j=1}^J a_{ij}^2,$$

where a_{ij} is firm j ’s share of the region’s total patent applications. The index is formed by first summing applications by firm within each region-year and then taking the concentration of those shares — the construction matters, and is returned to in Section VI. Regional human capital is measured following the use of education data by Aghion et al. (2009): HC , the number of students enrolled at universities and technical colleges, expressed per head of the regional population. Finally, urbanisation is measured by $Urban$, the share of a region’s population living in municipalities of more than ten thousand inhabitants, a control motivated by the work of Brenner and Greif (2006) on agglomeration and innovation.

V Empirical strategy

The dependent variable is a count: the number of patent applications published by a firm in a given year. Counts are non-negative integers, and the natural baseline for modelling them is the Poisson regression, in which the conditional mean is $E[y_{it} | x_{it}] = \exp(x'_{it}\beta)$ — a generalised linear model in the sense of Nelder and Wedderburn (1972). The exponential form keeps the predicted mean positive and gives the coefficients a convenient interpretation, discussed below.

Two features of the data shape the specification. The first is unobserved heterogeneity between firms. Firms differ in ways that are not captured by the regressors — in

their technological opportunities, their research traditions, the quality of their management — and these differences plausibly correlate both with size and with the propensity to patent. Ignoring them would confound the size–innovation relationship with fixed differences between large and small firms. The analysis therefore uses a fixed-effects specification, $E[y_{it} \mid x_{it}, \alpha_i] = \alpha_i \exp(x'_{it}\beta)$, in which the firm-specific term α_i absorbs all time-invariant firm characteristics. The fixed-effects Poisson estimator, introduced by Hausman et al. (1984) in their study of the patents–R&D relationship, is well suited to firm-level patent panels: unlike the fixed-effects logit estimator, it is consistent and free of the incidental-parameters problem, so that firm effects can be included without biasing the slope coefficients. A Hausman test decisively rejects the random-effects alternative ($\chi^2 = 53.4$, $p < 0.001$), and fixed effects are used throughout.

The second feature is that the Poisson model assumes equidispersion — a variance equal to the mean — which patent counts routinely violate, the variance typically exceeding the mean. Under this violation the conventional Poisson standard errors are not reliable. The estimates are therefore reported with standard errors clustered by firm, in the manner of Wooldridge (1999), which remain valid under both overdispersion and arbitrary correlation among the repeated observations of a firm (Cameron and Trivedi, 2015; Winkelmann, 2008).

The estimating equation can now be stated. For firm i in region j and year t , the expected number of published patent applications is

$$E[P_{it} \mid x_{it}, \alpha_i] = \alpha_i \exp \left(\beta_1 \ln S_{it} + \beta_2 (\ln S_{it})^2 + \beta_3 (\ln S_{it})^3 + \beta_4 \text{Age}_{it} \right. \\ \left. + \beta_5 \text{HHI}_{jt} + \beta_6 \text{HC}_{jt} + \beta_7 \text{Urban}_{jt} + \beta_8 D_{it}^{AG} \right),$$

where $\ln S$ is the logarithm of the firm’s size class, Age its age, HHI the regional patent-concentration index, HC regional human capital, Urban the degree of urbanisation, and D^{AG} a dummy for access to the capital market through the joint-stock form.

The central specification choice is the treatment of firm size. Size enters as a third-degree polynomial in $\ln S$. A single linear term would impose precisely the monotonic relationship that is under test — innovation rising steadily with size, as the Schumpeterian hypothesis asserts — and so would assume the answer rather than examine it. A cubic, by contrast, allows the estimated relationship to rise, fall and rise again, or to take any simpler shape the data prefer, and so can accommodate the non-monotonic, U-shaped pattern that earlier studies such as Pavitt et al. (1987) had found. The model is estimated for the full sample and then separately for process and for product innovations, and for the sub-periods before and after the Four-Year Plan of 1936, the latter split designed to test whether the shift to a rearmament regime altered the size–innovation relationship.

Because the dependent variable is a count, the estimation panel includes the firm-years in which a firm filed no application at all. Each firm is observed across its active span

— from its first to its last recorded application — and the intervening years without an application enter as genuine zeros; firm size in those years is carried forward from the nearest balance sheet. This yields a panel of 2,034 firm-years for 493 firm-locations, of which 584 are zeros. The construction departs from the underlying dissertation, which estimated on the firm-years with at least one application only: that restriction conditions the sample on a positive outcome and discards the quiet years that a count model is meant to use. Estimating on the restricted sample instead leaves the central result in place but, as the firm size variable shows below, somewhat understates it.

Finally, a note on interpretation. Because the conditional mean is exponential, $\partial E[y_{it} | x_{it}, \alpha_i] / \partial x_{j,it} = \beta_j \cdot E[y_{it} | x_{it}, \alpha_i]$. For a regressor entered in levels, such as firm age, β_j is a semi-elasticity — the proportional change in expected applications associated with a one-unit change in the regressor. For a regressor entered in logarithms, such as the size terms, β_j is an elasticity.

VI Results

Table 2 reports the fixed-effects Poisson estimates for the full sample, on the count panel described in Section V. The size polynomial carries the result. The linear term $\ln S$ is positive, the quadratic term $(\ln S)^2$ negative, and the cubic term $(\ln S)^3$ positive again; the three are jointly significant on a Wald test ($\chi^2(3) = 15.1, p = 0.002$). This alternation of signs is incompatible with a monotonic relationship between firm size and innovation: a straight-line Schumpeterian effect would require $\ln S$ alone to carry the relationship, with the higher-order terms indistinguishable from zero.

The shape behind these coefficients is shown by the solid curve in Figure 4, which plots the fitted conditional mean across the six size classes, normalised to the smallest class. Inventive activity is highest at the two ends of the size distribution — among the smallest firms and again among the great concerns — and lowest among medium and medium-large firms. The implication for the hypothesis under test is twofold. The monotonic version of Schumpeter’s second hypothesis, that innovativeness rises steadily with size, is rejected: the relationship is not monotonic but U-shaped, or more precisely cubic. Yet the result is not simply anti-Schumpeterian. The largest concerns are among the most inventive firms in the industry; what fails is only the claim that size as such, across the whole range, brings ever greater innovativeness. The pattern instead recalls the U-shaped size distribution of innovating firms that Pavitt et al. (1987) found for post-war Britain, in which both the small and the very large innovate intensively while the middle of the distribution lags.

To formalise this non-monotonic shape, a Sasabuchi-style test (Lind and Mehlum, 2010), extended to the cubic, asks three questions, each using the cluster-robust covariance of the size coefficients. First, is the slope of the fitted relationship positive at the lower

Table 2: Fixed-effects Poisson estimates, full sample. Dependent variable: published patent applications per firm-year.

	Coefficient
$\ln S$	2.48** (0.80)
$(\ln S)^2$	-3.22** (1.24)
$(\ln S)^3$	1.15* (0.47)
Age	-0.007 (0.011)
Patent concentration (HHI)	0.29 (0.31)
Human capital	13.97 (31.39)
Urban	-2.93 (2.08)
AG dummy	0.04 (0.18)
Firm-years N	2,034
of which zeros	584
Firm-locations n	493
Wald $\chi^2(3)$, size terms	15.1
p -value	0.002

Notes: Conditional fixed-effects Poisson on the count panel including zero-application firm-years; cluster-robust standard errors (by firm) in parentheses. *, **, *** denote significance at the 5, 1 and 0.1 per cent level. Patent concentration is the regional Herfindahl of firms' patent shares (see text). The Wald statistic tests the joint hypothesis $\ln S = (\ln S)^2 = (\ln S)^3 = 0$.

bound of the observed size range? At size class 1 ($\ln S = 0$) the slope is $\beta_1 = 2.48$ with standard error 0.80 ($t = 3.12$, one-sided $p = 0.0009$). Second, is the slope positive at the upper bound? At size class 6 ($\ln S = \log 6 \approx 1.79$) the slope is the linear combination $\beta_1 + 2\beta_2 \log 6 + 3\beta_3(\log 6)^2$, evaluated at 1.99 with standard error 0.88 ($t = 2.28$, one-sided $p = 0.011$). Third, does the cubic have two real interior turning points, so that the curve genuinely rises, falls, and rises again rather than running monotonically through the size range? The discriminant $\beta_2^2 - 3\beta_1\beta_3 = 1.81$ is positive, and the two turning points sit at size classes $S^* \approx 1.7$ and $S^* \approx 3.8$, both inside the observed range $[1, 6]$. The joint one-sided p -value of the two endpoint tests is 0.011, so the monotonic null — a single sign of the slope across the entire size range, or no interior turning point — is rejected at conventional levels. The data support the rising–falling–rising cubic shape unambiguously.

The control variables call for caution. Market structure is measured here by a single, correctly constructed index: the regional Herfindahl of firms' patent shares, formed by summing applications by firm within each region-year and then taking the concentra-

tion of those shares. This replaces the two market-structure variables of the underlying dissertation, which were both built as Herfindahl indices of quantities that are constant within a region-year and so collapsed towards the reciprocal of the local count of applications rather than measuring concentration across firms. The corrected index enters the full-sample regression positive but insignificant ($0.29, z \approx 0.9$). Its sign is in the Schumpeterian direction — concentration associated with more inventive activity — and, as the sub-samples below show, it becomes significant for process innovations and for the post-1936 years; but it is not robust across specifications, and the evidence on competition and concentration is best regarded as suggestive rather than settled. The degree of urbanisation, the AG dummy and firm age are small and statistically indistinguishable from zero.

The estimate on regional human capital requires a frank comment. The point estimate is large and positive, which on its face would suggest a substantial role for the supply of university- and technical-college-trained labour. It is, however, statistically insignificant: the cluster-robust standard error is about 31 against a coefficient of 14, so that $z \approx 0.4$. This matters because the underlying dissertation reported standard errors for this term of the order of 0.01 — implying t -statistics in the thousands — which were not genuine. In the fixed-effects specification the human-capital variable varies overwhelmingly between regions and very little within a given firm over time, so the model is weakly identified along that dimension; the near-singular information matrix was inverted in the original code with a generalised inverse, which collapses the estimated variance towards zero. A regular matrix inverse settles the question: the term is large but insignificant here and in every sub-sample below. The apparent “strong” association between regional human capital and patenting does not survive a correct calculation of the standard errors, and the result is best set aside pending the better-identified design discussed in Section VII.

The remaining results concern heterogeneity, and are reported in Table 3. The first split separates process from product innovations, and here the corrected estimation parts company with the dissertation. The cubic size pattern is sharp and jointly significant for process innovations ($\chi^2(3) = 37.0, p < 0.001$) but not for product innovations ($\chi^2(3) = 5.5, p = 0.14$); the dissertation’s specification, estimated on the restricted sample, showed if anything the reverse. The honest reading is that the split between innovation types is not robust: the non-monotonic relationship is a feature of the pooled data and of the process-innovation subsample, while for product innovations alone it cannot be established with confidence. Little should be built on the contrast between the two types.

The second split is the one that bears most directly on the historical argument: the years before and after the Four-Year Plan of 1936. Columns 3 and 4 of Table 3 estimate the model separately for the two sub-periods, and the two further curves in Figure 4 show the corresponding fitted shapes. The contrast is stark. Before 1936 the cubic size pattern is sharp and overwhelmingly significant ($\chi^2(3) = 92.6, p < 0.001$): the curve

Table 3: Fixed-effects Poisson estimates by innovation type and sub-period. Count panel including zero-application firm-years.

	Process	Product	Before 1936	After 1936
$\ln S$	4.89** (1.63)	1.51 (0.92)	3.95*** (0.50)	1.21 (2.08)
$(\ln S)^2$	-5.76* (2.78)	-2.18 (1.44)	-5.80*** (0.83)	-1.67 (1.99)
$(\ln S)^3$	1.98 (1.07)	0.82 (0.54)	2.18*** (0.33)	0.61 (0.60)
Age	0.005 (0.016)	-0.007 (0.010)	0.005 (0.019)	-0.028 (0.026)
Patent conc. (HHI)	0.78* (0.38)	0.24 (0.37)	-0.20 (0.39)	0.66* (0.29)
Human capital	7.18 (28.43)	50.47 (39.49)	-29.10 (39.04)	3.54 (39.98)
Urban	-2.82 (2.15)	-1.85 (1.88)	-8.44 (4.53)	0.87 (5.22)
AG dummy	0.05 (0.39)	-0.11 (0.15)	0.36** (0.14)	-0.26 (0.58)
Firm-years N	1,047	1,587	965	784
of which zeros	302	466	189	111
Firm-locations n	289	373	318	304
Wald $\chi^2(3)$, size	37.0	5.5	92.6	1.8
p -value	< 0.001	0.14	< 0.001	0.62

Notes: Conditional fixed-effects Poisson on the count panel including zero-application firm-years; cluster-robust standard errors (by firm) in parentheses. *, **, *** denote significance at the 5, 1 and 0.1 per cent level. Patent concentration is the regional Herfindahl of firms' patent shares. The Wald statistic tests the joint hypothesis $\ln S = (\ln S)^2 = (\ln S)^3 = 0$.

climbs steeply towards the great concerns, the largest firms reaching a fitted inventive activity well above both the smallest firms and the mid-sized trough. After 1936 the size terms collapse — none is individually significant and the joint Wald test cannot reject the hypothesis that all three are zero ($\chi^2(3) = 1.8$, $p = 0.62$). On the face of it, then, the non-monotonic size pattern is a feature of the pre-rearmament years and disappears thereafter; the rearmament years did not concentrate inventive activity on the great concerns. This is not the reading the patent series can ultimately support, however. As Section VII sets out, the published-application measure becomes progressively less reliable after 1936, precisely because the inventions of the large armament concerns were the ones most likely to be patented in secret and so to vanish from the published record. The post-1936 result is therefore better treated as a limitation of the source than as a finding about the war economy: the rearmament period is the stretch over which these data can say least.

A natural further check separates the heavy upstream of the panel — mining, iron and steel, the firms closest to the coal-and-steel core of the historical Montanindustrie — from the rest. The upstream sub-set accounts for 784 of the 2,034 firm-years and 181 of the




Figure 4: Fitted size–innovation relationship across the six firm size classes (conditional mean of published patent applications, normalised to class 1), for the full sample and the two sub-periods. Shaded regions are 95 per cent confidence bands built from the cluster-robust covariance of the size terms $(\ln S, (\ln S)^2, (\ln S)^3)$. The bands pinch at class 1 by construction and widen at the upper end of the size distribution, where the fitted curve depends most heavily on the cubic term; the After-1936 band is wide throughout because the size terms cannot be jointly distinguished from zero in that sub-period (see Table 3).

493 firm-locations. Figure 5 shows that the curve estimated on this sub-set is essentially the full-sample curve, only slightly steeper: the same peak at the smaller end of the size distribution, the same trough in the middle, and the same recovery towards the great concerns. A finer split by downstream sector could not be reported reliably, because firms outside the upstream of the industry show too little within-firm variation in size class over the period 1928–1941 to identify the cubic under fixed effects — the linear-term standard error is of order nine when the model is estimated on machinery firms alone. The non-monotonic pattern is therefore best read as a feature of the panel as a whole and of its upstream heavy-industry core, rather than as the artefact of a single downstream sector.

Robustness checks. Two further specifications probe whether the cubic shape of the size–innovation relationship survives departures from the baseline conditional fixed-effects Poisson. The first adds year fixed effects on top of the firm fixed effects, absorbing any common annual shock to publication intensity — the diffusion of new patent classes, the state of the business cycle, the secrecy regime that tightened after 1936. The second estimates the same model as a pooled negative binomial with year fixed effects, the standard alternative count specification for overdispersed data (Cameron and Trivedi, 2015); the firm fixed effects are dropped, because the conditional fixed-effects negative-binomial estimator of Hausman et al. (1984) does not in fact eliminate the firm-specific dispersion




Figure 5: Sector heterogeneity: fitted size–innovation curve for the upstream mining, iron and steel sub-set of the panel, with the full-sample curve as a reference. Shaded regions are 95 per cent confidence bands from the cluster-robust covariance of the size terms $(\ln S, (\ln S)^2, (\ln S)^3)$. The two curves track one another closely; the upstream pattern is essentially the pooled pattern of Figure 4, only slightly more pronounced.

parameter (Allison and Waterman, 2002; Greene, 2008), and the comparable parameterisation is the pooled negative binomial with cluster-robust standard errors. Table 4 reports the size coefficients across the three specifications.

Taken together, the results give a clear answer to the question posed in Section III. In the German coal and steel industry between 1928 and 1941, large firms were not uniformly more innovative than small ones: the relationship between size and inventive activity was non-monotonic, with the smallest firms and the great concerns the most inventive and the middle of the size distribution the least. The monotonic form of Schumpeter’s second hypothesis is rejected. The evidence on market structure is suggestive but not robust: the regional patent-concentration index enters positive — in the Schumpeterian direction — and is significant for process innovations and for the post-1936 years, but not for the full sample. The non-monotonic size pattern is sharp before 1936 and cannot be detected after it; the published-patent record shows the size gradient dissolving in the rearmament years, although that stretch of the record is the least trustworthy. And the apparent role of regional human capital, once its standard errors are computed correctly, cannot be sustained: the coefficient is large but statistically insignificant.

Table 4: Robustness: the size cubic across alternative specifications.

	Baseline FE Poisson	+ Year FE FE Poisson	Pooled NB + Year FE
$\ln S$	2.48** (0.80)	3.06** (1.09)	0.18 (0.44)
$(\ln S)^2$	-3.22** (1.24)	-3.83 (†)	-0.44 (0.60)
$(\ln S)^3$	1.15* (0.47)	1.31 (†)	0.41 (0.22)
Firm fixed effects	yes	yes	no
Year fixed effects	no	yes	yes
Firm-years N	2,034	2,034	2,034
Wald $\chi^2(3)$, size	15.1	—	—

Notes: Cluster-robust standard errors (by firm) in parentheses where defined. *, **, *** denote significance at the 5, 1 and 0.1 per cent level. † In the year-fixed-effects column the cluster-robust standard errors for the higher-order size terms could not be reported because of near-collinearity

between the year dummies and the time-varying regional regressors (*HHI*, *HC*, *Urban*): the conditional-likelihood Hessian remains positive definite and the point estimates are well identified, but the sandwich variance has numerically negative diagonal entries for $(\ln S)^2$ and $(\ln S)^3$. The cubic pattern visibly survives: the linear, quadratic and cubic size terms keep the same signs and similar magnitudes as in the baseline. The pooled negative binomial, which uses both between- and within-firm variation in size, returns smaller coefficients of the same signs; the conditional-mean shape is the same rising–falling–rising cubic, only attenuated because the between-firm component of the size–patenting association is weaker than the within-firm component.

VII Discussion and limitations

The central finding of this study is a non-monotonic relationship between firm size and inventive activity. It is worth being precise about what this does and does not imply for the debate set out in Section III. It refutes the strong, monotonic reading of Schumpeter’s second hypothesis — the claim that innovativeness rises steadily with size across the whole distribution. But it does not vindicate the opposite, small-firm view either. The great concerns of the German coal and steel industry were among its most inventive firms; what the data deny is only that size as such, at every margin, brought more innovation. The pattern is better read through the institutional lens of Section III than as a verdict for one side of the size debate. Inventive activity in the industry had two distinct centres of gravity: small, often specialised firms at one end, and the large, vertically integrated concerns at the other, with medium and medium-large firms — large enough to have lost the focus of the small specialist, not yet large enough to command the research resources of the great concern — the least inventive of all. The “innovative few” of Section IV and the U-shape of Section VI describe the same underlying structure.

The second substantive result, the behaviour of the size gradient around 1936, calls for more caution still. The naive reading of the sub-period estimates is striking: the

size–innovation curve is steeper *before* the Four-Year Plan than after it, so that on the published-patent measure the great concerns’ inventive lead was widest in the Depression and early-recovery years and narrowed once rearmament was under way. Taken literally, this would stand the conventional picture of the Nazi war economy on its head. But it should not be taken literally. The published-application series is exactly the wrong instrument for the rearmament years, because from 1939 the inventions of the large armament concerns — the firms at the top of the size distribution — were the ones most likely to be routed into secret patents and so to disappear from the record (see below). A measure that increasingly omits the war-relevant inventions of the biggest firms will mechanically flatten the size gradient after 1936. The honest conclusion is therefore a negative one: these data cannot establish how rearmament reshaped the size–innovation relationship, because the period of interest is also the period over which the measure is least complete. The pre-1936 estimates, by contrast, fall almost entirely within the years before secret patenting, and the strong U-shape they show is reliable.

These cautions point to the broader limitation of the analysis: it is associational, not causal. The fixed-effects specification removes all time-invariant differences between firms, which is a considerable advance over a cross-sectional comparison, but it cannot remove time-varying confounding, and none of the regressors is randomly assigned. Firm size is plainly endogenous to inventive success over a longer horizon — innovative firms grow — and the regional variables are endogenous in a subtler way: human capital, competition and inventive activity are all outcomes of the same regional economic development, so that the coefficients on regional patent concentration and on regional human capital should be read as conditional correlations, not as the effects of interventions. The human-capital estimates carry the further, specific problem set out in Section VI. In a fixed-effects model a regional variable that barely moves within a firm over time is only weakly identified, and the very small standard errors in the source material were an artefact of that weak identification compounded by the use of a generalised inverse. Re-estimating the model with a regular inverse removes the artefact and yields a cluster-robust standard error of about thirty for the human-capital coefficient: the term, large as its point estimate is, is not statistically significant. The result reported in the underlying dissertation — a strong association between regional human capital and patenting — therefore does not hold, and the question is best left to a design that can identify a regional variable properly, such as one exploiting cross-regional rather than within-firm variation.

A limitation inherent in the measure itself deserves emphasis. Patents capture inventive activity imperfectly: not every invention is patented, the propensity to patent differs across firms, and it is known to rise with firm size, as large firms patent more strategically and more completely than small ones. In a study whose subject is precisely the relationship between size and patenting, this is not a minor caveat. Part of the size gradient estimated here may reflect differences in the propensity to patent rather than in inventive

activity, and the data cannot separate the two. The non-monotonic *shape*, however — the fact that medium-sized firms patent less than both smaller and larger ones — is not easily explained by a propensity-to-patent story, which would predict a monotonic size gradient.




Figure 6: Published patent applications per year, 1928–1943. The series for 1943 is incomplete: the last issue of *Stahl und Eisen* used here appeared on 20 May 1943.

Finally, the war years bias the measure downward, and increasingly so. Figure 6 shows the annual count of published applications: it drifts downward across the whole period and then falls away in the early 1940s. Part of this decline is real, but part is an artefact of wartime patent administration. From September 1939 inventions important to national defence could be patented in secret, without publication; the Reich increasingly took such applications over as trustee; and from 1943 the patent office was wound down, the last issue used here appearing on 20 May 1943. The post-1936 sub-period is therefore measured with a downward bias that worsens as the war goes on, and it is the great armament concerns whose inventions were most likely to be withdrawn into secrecy. This is what makes the post-1936 estimates so hard to read: because secrecy removes disproportionately the patents of the largest firms, it mechanically flattens the measured size gradient, and the flatter post-1936 curve of Section VI is most plausibly an artefact of exactly this process rather than evidence of a real change in the inventiveness of the great concerns. Taken together, these limitations bear unevenly on the findings. The non-monotonic size-innovation relationship is robust, and it is most securely estimated in the pre-1936 years, before secret patenting set in; the behaviour of the size gradient during the rearmament period cannot be established from the published-patent record at all; and the apparent role of regional human capital, once its standard errors are computed correctly, does not hold. What survives, and survives strongly, is the central result: that firm size and inventive activity were related non-monotonically, and that the medium-sized firm, not

the small one, was the laggard of the industry.

VIII Conclusion

This paper has used a newly assembled firm-level dataset to test one of the oldest propositions in the economics of innovation — Schumpeter’s claim that large firms are the engines of technical progress — in a setting that has not been examined in this way before: the German coal and steel industry between 1928 and 1941. The dataset draws together more than ten thousand published patent applications, firm balance-sheet data from the *Handbücher der deutschen Aktiengesellschaften*, and a range of regional statistics, and it allows the size–innovation relationship to be estimated with firm fixed effects across a turbulent decade and a half.

The answer it returns is a qualified one. In this industry, over this period, large firms were not uniformly more innovative than small ones. The relationship between firm size and inventive activity was non-monotonic: the smallest firms and the great concerns were the most inventive, and the middle of the size distribution the least. The strong, monotonic form of Schumpeter’s second hypothesis is rejected. But the result is not a simple refutation of Schumpeter, for the largest concerns were indeed among the most inventive firms in the industry; what fails is the claim that size brings ever greater innovativeness at every margin. The market-structure evidence is suggestive but not settled: a correctly constructed Herfindahl of regional patent concentration enters in the Schumpeterian direction — positive — and is significant for process innovations and for the post-1936 years, but not for the full sample, so the analysis cannot firmly adjudicate between the competition and concentration views. The size–innovation relationship was estimated separately for the years before and after the Four-Year Plan of 1936; it is sharp and strongly significant before 1936 and cannot be detected after it, but the spread of secret patenting in the war years makes that later period too poorly measured to interpret, so the firmest sub-period evidence is the strong U-shape of the pre-1936 years.

These findings speak to two literatures. For the long debate on firm size and innovation, the German coal and steel industry adds a case in which the relationship is neither the monotonic one of the Schumpeterian tradition nor the small-firm one of its critics, but the U-shape that a handful of earlier studies had glimpsed — evidence that the shape of the relationship is contingent on industry and institutional setting rather than universal. For the economic history of the Nazi war economy, the results are more sobering than affirmative: they show how quickly the published-patent record ceases to be usable once secret patenting takes hold, and so how cautiously inventive activity under the war economy must be quantified at all.

Three lines of further work follow directly. The first concerns the rearmament period: because the published series cannot speak to it, a serious treatment of innovation under

the war economy would have to draw on the records of secret patents and on firm archives rather than on *Stahl und Eisen* alone. The second extends the analysis from firm size to the other half of the regional innovation system: the same dataset records the cooperation between firms and their specialisation across patent classes and regions, and a companion paper will examine whether regional knowledge spillovers and inter-firm cooperation shaped innovation alongside size. The third is the construction of the data itself — extending the firm panel beyond the patenting firms, and linking it to employment and output, would allow the propensity-to-patent problem to be addressed directly, and would give regional human capital the better-identified test that the fixed-effects design here could not provide. The broader contribution, however, is already clear: that the relationship between firm size and innovation, far from being a universal law, took a specific and historically contingent shape in one of the central industries of the inter-war German economy.

Appendix A. Data and variable definitions

A.1 Sources

The patent record is the published-application series of the German Reich Patent Office (*Reichspatentamt*), retrieved by firm name and inventor residence for the years 1928–1943 and back-dated to the year of application (eighteen months before publication, hence the suffix `_18M` on time-varying variables). Firm-level financial and organisational variables — equity capital, joint-stock status, age at founding, location — come from the annual *Handbücher der deutschen Aktiengesellschaften* for the period (Handbuch der deutschen Aktiengesellschaften, 1933–1944), the standard reference for German company finance in the inter-war years. Sector classifications, regional populations and the regional totals of university and technical-college students are taken from the Reich statistical yearbooks, the HISTAT historical statistics archive, and Titze’s data handbook of German educational history (Titze, 1987; Grupp et al., 2002).

A.2 Variables used in the regression

Table 5 lists every variable that enters the conditional fixed-effects Poisson estimation of Tables 2 and 3.

A.3 Sample construction

The estimation panel is built in four steps.

Step 1: record level. The source data are at the level of an individual published patent application linked to one or two named firms. After dropping rows whose first-named firm

is foreign (*Ausland_Firma1*=1) and the application years 1926–1927 (back-dated from 1928 publications, used only for back-dating), and after appending second-named domestic or Austrian firms as additional firm-record observations, this layer contains roughly ten thousand application–firm records.

Step 2: firm-year aggregation. Records are aggregated to the firm-location level, summing applications across patents within each firm-location and year. A firm-location is the unique combination of firm name (*FirmaKurz*) and regional district (*Regierungsbezirk*); firms operating in more than one district enter as separate firm-location observations, as in the underlying dissertation. After aggregation the firm-year panel contains 1,450 firm-years with at least one application, for 484 firm-locations.

Step 3: extension to the count panel with zeros. Because the dependent variable is a count, the estimation panel should include the firm-years in which a firm filed no application at all. Each firm-location is therefore observed across its active span — from the first to the last calendar year in which it filed at least one application — and the intervening application-free years enter as genuine zeros. Firm size and the AG dummy are carried forward within the firm-location over those zero years; age is incremented by one per calendar year; the regional variables are merged in afresh by region-year and carried forward within region where missing. The result is the panel of 2,034 firm-year observations on 493 firm-locations, of which 584 are zeros, used in Tables 2, 3 and Figures 4 and 5.

Step 4: complete-case restriction. Observations missing any regression variable after the carry-forward step are dropped. The dominant cause of attrition is missing equity capital in some firm-years (zero or unrecorded), which translates into a missing size class; such firm-years are excluded from the estimation panel rather than being folded into the smallest size class, as the underlying dissertation did.

A.4 Computational notes

All estimates in this paper are obtained from the conditional fixed-effects Poisson estimator of Hausman et al. (1984), equivalent to a multinomial logit of within-firm patent shares conditional on the firm’s total. Cluster-robust standard errors are constructed by the sandwich $(-H)^{-1}(\sum_i g_i g_i')(-H)^{-1}$, where g_i is the firm- i contribution to the conditional score and H the conditional Hessian, both evaluated at the maximum-likelihood estimate. The conditional Hessian is inverted with an ordinary matrix inverse rather than with the generalised inverse used in the underlying dissertation; the difference matters most for variables that are weakly identified within firms (here regional human capital), whose standard errors collapsed towards zero under the generalised inverse but are correctly reported at the order of 25–40 here.

References

- Zoltan J. Acs and David B. Audretsch. Innovation, market structure, and firm size. *The Review of Economics and Statistics*, 69(4):567–574, 1987.
- Zoltan J. Acs and David B. Audretsch. Innovation in large and small firms: An empirical analysis. *The American Economic Review*, 78(4):678–690, 1988.
- Philippe Aghion, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. Competition and innovation: An inverted-U relationship. *The Quarterly Journal of Economics*, 120(2):701–728, 2005.
- Philippe Aghion, Leah Boustan, Caroline Hoxby, and Jerome Vandenbussche. The causal impact of education on economic growth: Evidence from U.S. *Brookings Papers on Economic Activity*, 1:1–73, 2009.
- Paul D. Allison and Richard P. Waterman. Fixed-effects negative binomial regression models. *Sociological Methodology*, 32(1):247–265, 2002.
- David B. Audretsch and Zoltan J. Acs. Innovation and size at the firm level. *Southern Economic Journal*, 57(3):739–744, 1991.
- Jörg Baten, Anna Spadavecchia, Jochen Streb, and Shuxi Yin. What made southwest German firms innovative around 1900? assessing the importance of intra- and inter-industry externalities. *Oxford Economic Papers*, 59(Supplement):i105–i126, 2007.
- Knut Borchardt. *Grundriß der deutschen Wirtschaftsgeschichte*. Vandenhoeck & Ruprecht, Göttingen, 1985.
- Thomas Brenner and Siegfried Greif. The dependence of innovativeness on the local firm population: An empirical study of German patents. *Industry and Innovation*, 13(1): 21–39, 2006.
- Walther Busse von Colbe. *Die Planung der Betriebsgröße*. Betriebswirtschaftlicher Verlag Gabler, Wiesbaden, 1964.
- A. Colin Cameron and Pravin K. Trivedi. Count panel data. In Badi H. Baltagi, editor, *The Oxford Handbook of Panel Data*, chapter 8. Oxford University Press, Oxford, 2015.
- Forrest Capie. Disintegration of the international economy between the wars. In Nicholas Crafts and Peter Fearon, editors, *The Great Depression of the 1930s: Lessons for Today*, pages 140–164. Oxford University Press, Oxford, 2013.
- William S. Comanor. Research and technical change in the pharmaceutical industry. *The Review of Economics and Statistics*, 47(2):182–190, 1965.

- William S. Comanor and Frederic M. Scherer. Patent statistics as a measure of technical change. *Journal of Political Economy*, 77(3):392–398, 1969.
- Harald Degner. *Sind große Unternehmen innovativ oder werden innovative Unternehmen groß? Eine Erklärung des unterschiedlichen Innovationspotenzials von Unternehmen und Regionen*. Jan Thorbecke Verlag, Ostfildern, 2012.
- Alexander Donges. *Die Vereinigte Stahlwerke AG im Nationalsozialismus: Konzernpolitik zwischen Marktwirtschaft und Staatswirtschaft*. Verlag Ferdinand Schöningh, Paderborn, 2014.
- Helmut Fiereder. *Reichswerke Hermann Göring in Österreich (1938–1945)*. Geyer-Edition, Wien and Salzburg, 1983.
- Armin Frisch. *Unternehmensgröße und Innovation: die schumpeterianische Diskussion und ihre Alternativen*. Campus-Verlag, Frankfurt am Main and New York, 1993.
- Richard Gilbert. Looking for Mr. Schumpeter: Where are we in the competition–innovation debate? In Adam B. Jaffe, Josh Lerner, and Scott Stern, editors, *Innovation Policy and the Economy, Volume 6*, pages 159–215. MIT Press, Cambridge, MA, 2006.
- William Greene. Functional forms for the negative binomial model for count data. *Economics Letters*, 99(3):585–590, 2008.
- Zvi Griliches. Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28(4):1661–1707, 1990.
- Hariolf Grupp. *Messung und Erklärung des technischen Wandels: Grundzüge einer empirischen Innovationsökonomik*. Springer, Berlin and Heidelberg, 1997.
- Hariolf Grupp, Iciar Dominguez-Lacasa, and Monika Friedrich-Nishio. *Das deutsche Innovationssystem seit der Reichsgründung: Indikatoren einer nationalen Wissenschafts- und Technikgeschichte in unterschiedlichen Regierungs- und Gebietsstrukturen*. Physica Verlag, Heidelberg, 2002.
- Handbuch der deutschen Aktiengesellschaften. *Handbuch der deutschen Aktiengesellschaften, Jahrgänge 38–49*. Verlag Hoppenstedt, Berlin, 1933–1944.
- Jerry Hausman, Bronwyn H. Hall, and Zvi Griliches. Econometric models for count data with an application to the patents–R&D relationship. *Econometrica*, 52(4):909–938, 1984.
- Morton I. Kamien and Nancy L. Schwartz. Market structure and innovation: A survey. *Journal of Economic Literature*, 13(1):1–37, 1975.

- Christopher Kopper. *Hjalmar Schacht: Aufstieg und Fall von Hitlers mächtigstem Bankier*. Carl Hanser Verlag, München and Wien, 2006.
- Jo Thori Lind and Halvor Mehlum. With or without U? the appropriate test for a U-shaped relationship. *Oxford Bulletin of Economics and Statistics*, 72(1):109–118, 2010.
- Edwin Mansfield. Size of firm, market structure, and innovation. *Journal of Political Economy*, 71(6):556–576, 1963.
- Edwin Mansfield. Patents and innovation: An empirical study. *Management Science*, 32(2):173–181, 1986.
- Gerhard Mollin. *Montankonzerne und Drittes Reich: Der Gegensatz zwischen Monopolindustrie und Befehlswirtschaft in der deutschen Rüstung und Expansion 1936–1944*. Vandenhoeck & Ruprecht, Göttingen, 1988.
- John A. Nelder and Robert W. M. Wedderburn. Generalized linear models. *Journal of the Royal Statistical Society. Series A (General)*, 135(3):370–384, 1972.
- Manfred Neumann, Ingo Böbel, and Alfred Haid. Innovations and market structure in west german industries. *Managerial and Decision Economics*, 3(3):131–139, 1982.
- Tom Nicholas. Why Schumpeter was right: Innovation, market power, and creative destruction in 1920s America. *The Journal of Economic History*, 63(4):1023–1058, 2003.
- Patrizio Pagano and Fabiano Schivardi. Firm size distribution and growth. *The Scandinavian Journal of Economics*, 105(2):255–274, 2003.
- Ariel Pakes and Zvi Griliches. Patents and R&D at the firm level: A first report. *Economics Letters*, 5(4):377–381, 1980.
- Keith Pavitt, Michael Robson, and Joe Townsend. The size distribution of innovating firms in the UK: 1945–1983. *The Journal of Industrial Economics*, 35(3):297–316, 1987.
- Almarin Phillips. Concentration, scale and technological change in selected manufacturing industries 1899–1939. *The Journal of Industrial Economics*, 4(3):179–193, 1956.
- Mark Rogers. Networks, firm size and innovation. *Small Business Economics*, 22(2):141–153, 2004.
- Joseph A. Schumpeter. *Theorie der wirtschaftlichen Entwicklung*. Kyklos, Wien, 1911.
- Joseph A. Schumpeter. *Capitalism, Socialism and Democracy*. Harper & Brothers, New York and London, 1942.

- Bruce Smith. Technological innovation in electric power generation: 1950–1970. *Land Economics*, 50(4):336–347, 1974.
- Donald Smythe. The great merger movement and the diffusion of electric power utilization in american manufacturing, 1899–1909. *Eastern Economic Journal*, 27(3):253–266, 2001.
- Werner Sombart. *Der moderne Kapitalismus, Band 3*. Duncker & Humblot, München and Leipzig, 1927.
- Jochen Streb, Jörg Baten, and Shuxi Yin. Technological and geographical knowledge spillover in the German empire, 1877–1918. *The Economic History Review*, 59(2):347–373, 2006.
- Catherine Sveikauskas and Leo Sveikauskas. Industry characteristics and productivity growth. *Southern Economic Journal*, 48(3):769–774, 1982.
- Hartmut Titze. *Datenhandbuch zur deutschen Bildungsgeschichte*. Vandenhoeck & Ruprecht, Göttingen, 1987.
- Andrea Vaona and Mario Pianta. Firm size and innovation in european manufacturing. *Small Business Economics*, 30(3):283–299, 2008.
- John M. Vernon and Peter Gusen. Technical change and firm size: The pharmaceutical industry. *The Review of Economics and Statistics*, 56(3):294–302, 1974.
- Rainer Winkelmann. *Econometric Analysis of Count Data*. Springer, Berlin and Heidelberg, 5 edition, 2008.
- Jeffrey M. Wooldridge. Distribution-free estimation of some nonlinear panel data models. *Journal of Econometrics*, 90(1):77–97, 1999.

Table 5: Variable definitions.

Name	Level	Definition
P_{it} (<i>Patents</i>)	firm-year	Count of patent applications published by firm i in year t , back-dated to the year of filing. Dependent variable. Sums over both first-named and second-named firms where the second-named firm is domestic or Austrian.
$\ln S_{it}$	firm-year	log of the firm's monotone equity-size class (1 = micro, 6 = great concern). Classes are bounded at ≤ 0.5 , 2, 5, 50 and 100 million Reichsmark of equity capital (<i>Eigenkapital</i>), with class 6 for firms above 100 million RM. Firm-years with zero or missing equity are treated as missing on this variable.
$(\ln S_{it})^2, (\ln S_{it})^3$	firm-year	Squared and cubed values of $\ln S_{it}$, included to admit a non-monotonic size-innovation relationship.
Age_{it}	firm-year	Years since founding, computed from the founding year listed in the <i>Handbuch</i> ; mergers count as new foundations. Carried forward across zero-application firm-years by +1 per calendar year.
HHI_{jt} (<i>Konk_Br</i>)	region-year	Regional Herfindahl-Hirschman index of patent shares, $\sum_f a_{fjt}^2$, where a_{fjt} is firm f 's share of total published applications by firms located in <i>Regierungsbezirk</i> j in year t . Computed by first summing applications by firm within each region-year and then taking the concentration of the resulting shares. Replaces the two mis-constructed concentration indices of the underlying dissertation, both of which collapsed to the reciprocal of the local count of applications.
HC_{jt}	region-year	Regional human capital: students enrolled at universities and technical colleges in region j at t , expressed per head of regional population.
$Urban_{jt}$	region-year	Share of the region's population living in municipalities of more than ten thousand inhabitants.
D_{it}^{AG} (<i>AG dummy</i>)	(AG firm-year	Indicator equal to one if firm i is organised as a joint-stock company (<i>Aktiengesellschaft</i>) in year t , zero otherwise.