

Clearance time and effective patent term in medical devices: evidence from EU and US records

Andreas Peters

a.peters81@icloud.com

Rosenheim University of Applied Sciences: Technische Hochschule Rosenheim

<https://orcid.org/0009-0006-0059-1420>

Research Article

Keywords: clearance time, effective patent term, notified bodies, medical devices, small and medium-sized enterprises, appropriability

Posted Date: October 16th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11021405/v1>

License:  This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Clearance time and effective patent term in medical devices: evidence from EU and US records

Andreas Peters

Rosenheim Technical University of Applied Sciences (TH Rosenheim), Rosenheim, Germany
ORCID 0009-0006-0059-1420 · a.peters81@icloud.com

Abstract

Analyses of regulatory friction in medical-device innovation stop at average clearance time. Patent term runs from filing and is consumed while a dossier waits, so appropriability depends on the whole distribution of clearance time. The author measures that distribution under the EU Medical Devices Regulation and its US counterparts from file-level Food and Drug Administration records, the notified-body survey and firm-size surveys. EU certification takes 18 to 22 months on average against 5.5 for a 510(k) clearance and 14.7 for a premarket approval, so the European route consumes about one to 1.4 years more patent term than the 510(k) route; US patents that obtain statutory term restoration for a device recover a median 2.62 years of that period, and no European counterpart exists. US records identify the tails exactly, the European record only from below; the ninetieth percentile of typical durations per notified body lies above 24 months without an upper bound, and 10 to 41 percent of bodies report durations outlasting the 23-month median financing interval, against under 1 percent of 510(k) files. A two-market valuation shows that US-first entry pays behind a 510(k) review only where completing the European dossier saves two to five and a half months, an amount no public record measures. Smaller manufacturers wait hardly longer on average, but their review phase has a longer tail; on the US record applicants filing once in five years wait 42 percent longer at the ninth decile than frequent applicants, holding device area and submission type fixed.

Keywords: clearance time; effective patent term; notified bodies; medical devices; small and medium-sized enterprises; appropriability

JEL Classification: O34; L51; I18; D81

1 INTRODUCTION

Administrative procedures shape industry dynamics through two prices, the money a firm pays and the time it waits, and the second is the harder of the two. Across 22 European countries and 14 manufacturing sectors, the time needed to complete bureaucratic procedures impedes entry and exit more than the monetary cost of the same procedures, and the effect falls hardest on smaller firms (Tomasi et al., 2023). Business dynamism suffers where administrative burden is high, because entry costs rise, experimentation falls and inefficient incumbents are shielded from displacement (Djankov et al., 2006; Calvino et al., 2020). Time enters that account as one component of the cost of doing business, recoverable in principle once the procedure ends.

Where the procedure is a market-access requirement rather than a registration formality, time enters differently. A medical device cannot be placed on the market before a notified body has certified it, while the patent protecting the device runs throughout, so every month spent in assessment is subtracted from the period in which the developer can earn a return on the invention. Administrative time here does not merely raise the cost of entry; it consumes the asset the firm entered to exploit, and the consumption is irreversible. Medical-device innovation in Europe is widely held to have slowed since the 2017 Medical Devices Regulation (MDR), and the usual diagnosis is that conformity assessment takes too long. Average duration captures only half of what that clock does. What the subtraction costs depends on how predictable clear-

ance is as well as on how long it takes, because firms able to plan against a known interval can finance, sequence and value projects that firms facing a wide and uncertain interval cannot, even at the same average (Online Resource 1, Fig. H.5).

In this article the author contributes to the literature on administrative burden and firm dynamics by measuring what a market-access clock does to appropriability, and by asking how that effect is distributed across firm types. Stated as a research question: how much appropriable patent term does regulatory clearance consume for medical devices under the MDR, compared with the US routes, once the whole distribution of clearance time and not only its mean is taken into account, and which firms bear the loss? Two quantities carry the answer. Effective patent term is the statutory term of twenty years minus the interval between the term-bearing filing and regulatory submission and minus clearance time, so it is what remains of protection when the product reaches the market. Clearance-time dispersion is the spread of that clearance interval around its mean, which determines not the average loss but its variability, and therefore what a firm must hold in reserve against it.

Four public records support the measurement. First, the notified-body survey commissioned by the European Commission reports, for each responding body, the typical time from a signed agreement to a certificate, in brackets (Austrian National Public Health Institute et al., 2026b). Second, MedTech Europe surveys manufacturers on certification fees and on phase durations separately for small and medium-sized enterprises and for large firms (MedTech Europe, 2024,

2022). Third, the US Food and Drug Administration (FDA) publishes the receipt and decision date of every premarket notification, De Novo request and premarket approval application, from which the clearance time of each dossier follows directly (Food and Drug Administration, 2026a,b). Fourth, Federal Register notices determining regulatory review periods for patent term restoration state, for each device, the dates of the investigational exemption, of the application and of approval (Federal Register, 2026); they establish that a restoration mechanism operates for US Class III devices, and they measure the phases it runs on, without revealing how much term is finally restored. Costs of capital, valuation multiples, financing intervals and turnover by firm size come from Damodaran, Carta and Eurostat (Damodaran, 2026; Carta, 2025, 2026; Eurostat, 2026).

Records of that kind identify different things, and the estimation strategy follows that asymmetry rather than concealing it. US durations are observed file by file, so means, quantiles and expected shortfalls are computed directly and need no distributional assumption. European durations are reported only as shares of bodies or of firms falling into duration brackets, so any statistic beyond a bracket boundary rests on a fitted distribution. The author therefore reports distribution-free bounds alongside every fitted European tail quantity, and validates the fitting procedure where validation is possible, by grouping the file-level 510(k) records into brackets, refitting, and comparing the fit with the truth. The exercise recovers the mean within about two percent and overstates the CV by about sixteen percent, which is why the European dispersion estimate is reported as an upper bound on precision rather than as a measurement. Two further threats are named where they arise. The European survey unit is the body and not the file, and the classification of applicants counts decisions across the whole window rather than only those preceding the submission at hand, which is why the estimation is repeated with a strictly backward-looking count and with applicant fixed effects. Where the record is rich enough, the firm-level question is estimated rather than described: the FDA file names the applicant of every premarket notification, so the number of submissions an applicant made in the observation window classifies applicants by how often they appear in the register, and the author regresses log clearance time on that classification at the mean and at the quartiles and the ninth decile, over the 15,702 of the 15,706 decisions whose recorded duration is positive, with submission type, third-party review, applicant country, fiscal year and device-area fixed effects, and standard errors clustered by applicant.

Measured this way, the European route consumes more patent term than either US comparator, which confirms in a market-access setting what the literature on administrative burden has found for registration procedures (Tomasi et al., 2023; Ciriaci, 2014).

Along three dimensions the author extends previous analyses. First, the loss is stated in the units in which protection is granted rather than in entry and exit rates. EU certification takes 18 to 22 months on average against 5.5 months for a 510(k) clearance and 14.7 months for an original premarket approval, which for a device that reaches review with 11.1 years of statutory term left consumes about one to 1.4 years more of it than the 510(k) route. Second, a two-market valua-

tion turns the timing question into a decision rule rather than a comparative statement. Sequencing entry so that the United States comes first is never optimal for a firm that would file the same European dossier anyway, and with observed costs, discount rates and financing terms it pays behind a 510(k) review only where completing the European file during that review saves between two and five and a half months, and behind a premarket approval only where it saves between six and fourteen. Third, the firm-size result is decomposed rather than asserted. Smaller manufacturers do not wait appreciably longer on average, but half of them report technical-documentation reviews above fifteen months against a quarter of large manufacturers, the fitted ninetieth percentiles stand at 34.3 against 19.5 months, and notified-body fees equal 0.28 years of turnover for the average enterprise with 10 to 19 persons employed against 0.02 years for one with 50 to 249. Pooling conceals the first of these at every level in the upper tail, and the mixture identity shows why: the between-group term accounts for 7.6 percent of pooled variance, so the burden sits in the within-group tail that pooled reporting averages away. The US record then tests the pattern at file level. Among 15,702 premarket notifications, applicants that submitted only once in five years wait 34 percent longer than applicants with twenty or more submissions at the median and 42 percent longer at the ninth decile, with submission type, device area and fiscal year held fixed, and their clearance distribution lies above that of frequent applicants at every quantile of practical interest. A strictly backward-looking count of prior decisions reproduces the gradient, so it does not rest on counting decisions the applicant had not yet obtained; a within-applicant specification on receipt cohorts leaves a residual gradient about a tenth as large, which places most of the difference between firms rather than inside them.

Closest to this one in question and in method is the study by Tomasi et al. (2023), which separates the cost and the time of administrative compliance and finds the time more damaging, particularly for small firms. The author's work differs in three respects. Their outcome is business churn and the author's is an appropriability asset with a fixed expiry, so time is not only a cost of entering but a subtraction from what entry is worth. Burden is measured there by a country-level index and here by the duration distribution of the procedure itself, which makes the dispersion of that duration, and not merely its level, an object of measurement. And their identification rests on differential sectoral exposure, whereas the author's rests on the contrast between a jurisdiction that publishes file-level durations and one that publishes brackets, which turns the reporting rule into an object of policy rather than a nuisance. Companion work by the author has modelled the mean-delay channel and the case for restoring term lost to clearance in two markets (Peters, 2026a), measured post-authorisation exclusivity for US devices from legal-status data (Peters, 2026b) and examined the legal design of a device-specific protection certificate (Peters, forthcoming); the dispersion channel and its distributional incidence are not treated there.

2 CLEARANCE TIME, APPROPRIABILITY AND FIRM SIZE

Standalone devices make the cleanest case for isolating the mechanism, and the two jurisdictions treat the loss very differently. In Europe no restoration instrument intervenes, since certification under the device regime does not satisfy the marketing-authorisation requirement of the regulation on supplementary protection certificates (Court of Justice of the European Union, 2018; European Parliament and Council, 2009). Whatever time the clearance process consumes is therefore lost to the innovator, which is why the spread of that time translates directly into the spread of the term that remains. In the United States an instrument exists for the highest-risk class. Under 35 U.S.C. §156 a patent claiming a device that passed premarket approval may be extended by part of the regulatory review period, subject to statutory deductions and to ceilings of five years and of fourteen years of remaining term after approval. The Food and Drug Administration determines that review period and publishes it; the median testing phase is 5.01 years in the sixty notices that state both phase dates and the median approval phase 0.86 years across all seventy (Federal Register, 2026). What the Patent and Trademark Office finally grants is published separately. Across the 161 device notices issued between 1997 and 2026, 223 distinct patents are named, and 111 of them appear on the office’s list of extensions granted under §156 (United States Patent and Trademark Office, 2026). Among those 111 the median grant is 956 days, or 2.62 years, with an interquartile range of 576 to 1,517 days and 20 percent at the five-year ceiling. Patents rather than devices are counted there, since a product carries at most one extended patent and the patents that obtained none lie outside both statistics. How much term the Patent and Trademark Office finally restores depends on deductions that those notices do not report, so the size of the recovery is not measured here; what is established is that a recovery mechanism exists on one side of the Atlantic and not on the other. Fig. 1 shows how much term each route consumes before any restoration, at the mean and at the ninetieth percentile, so the comparisons below understate the transatlantic difference. Pharmaceuticals, by contrast, recover part of the loss through a certificate, and combination products fall between the two regimes; the account below applies to the first case and states where it stops.

Appropriability research treats patent strategy as contingent on the institutional conditions under which value is captured, including standards, certification structures and the rules that govern access to complementary assets (Teece, 1986; Somaya, 2012; Holgersson et al., 2018). Firm-level evidence reads intellectual-property decisions in the same way, whether the object is the renewal of a trademark once portfolio composition makes it redundant (Agiropoulos et al., 2026), the rush to patent that accompanies a technological wave (Benassi et al., 2020), or the choices firms make about protection and disclosure inside university collaborations (Corsino and Torrisi, 2023); in each case the value of a right depends on conditions the firm does not set. Finite patent terms depress investment in innovations with long development periods (Budish et al., 2015), small firms respond more strongly to appropriability conditions than large incumbents (Holgersson, 2013), and




Fig. 1 Patent term consumed by clearance. Solid bars give the mean, hatched extensions the further wait up to the ninetyth percentile. US values are computed file by file from the FDA databases. The row for notified bodies uses the log-normal fit to the twentieth survey for the mean; its ninetyth percentile is bounded below by 24 months and not bounded above, because the reported counts place that percentile in the open upper bracket. The row for firms is the average technical-documentation certification time reported by MedTech Europe (2024), for which no distribution is published. What survives to market entry additionally depends on the date from which the term-bearing patent runs, which is not established here

shocks to appropriability travel asymmetrically by firm size (Aydin Ozden and Khashabi, 2023).

Industrial-organisation research on patent protection supplies the welfare side of the same question. Capuano and Grassi (2019) show that the level of protection decides whether an innovation is licensed, who pays most for it and how the market is structured afterwards, and Bennato and Giulietti (2019) extend the optimal-protection framework of Grossman and Lai (2004) to the externality that patent enforcement exerts on the availability of pharmaceutical products. Both treat the strength of protection as the policy variable; here the regulatory clock that consumes protection is a second one, because its dispersion changes how much of the statutory term is effectively available.

Regulation enters this literature mostly through its level or its expected duration. Peltzman (1973) measured the innovation cost of tighter drug approval; Stern (2017) documented a penalty for first entrants into new device categories when approval standards are uncertain; Grennan and Town (2020) used the difference between European and US approval standards for coronary stents to weigh the information that testing generates against the access it costs. Reviews of regulatory review separate its speed, its cost and its uncertainty (Boakye et al., 2022), patent protection dampens the negative effect of market uncertainty on R&D (Czarnitzki and Toole, 2011), and uncertainty about economic policy reduces patenting in a way consistent with real options, more strongly in high-technology industries and less in firms that preferential financing shields (Guoru et al., 2026; Chin et al., 2026). Read together, these results locate the transmission of regulatory friction in the timing of market entry and in the financing conditions of the firm, which is where the dispersion of a clearance clock operates; none of them maps that dispersion into the distribution of residual patent term, and none asks who bears it.

Sectoral-systems research adds that regulation belongs to the structure of a sector rather than constraining it from outside (Malerba, 2002), and mission-oriented innovation policy

directs attention to the capabilities of the public organisations that implement such structures (Mazzucato, 2018; Kattel and Mazzucato, 2018). Notified bodies are an unusual case, since private organisations exercise a public certification function under designation, and the capacity they build or fail to build is itself a policy outcome. They are also designated agents with finite capacity: above the lowest risk class a manufacturer cannot certify its own device, and excess demand is cleared by queues rather than by prices, so the conformity-assessment system behaves as a capacity-constrained complementor on which every innovator in the regulated segment depends (Granstrand and Holgersson, 2020; Granstrand, 2024). Congestion-pricing logic has been applied to backlogs at patent offices (Marco and Prieger, 2009); the present setting extends it to a queue in which patent term decays while the applicant waits.

Evidence on the MDR itself accumulated quickly after its date of application, documenting compliance difficulties, a differential burden on small and medium-sized enterprises (SMEs) and safety gains that remain difficult to demonstrate (Bretthauer et al., 2023; Huusko et al., 2023; Carl and Hochmann, 2024; Svempe, 2024; McDermott and Kearney, 2025). Documentation of this kind establishes that the transition was costly; it does not formalise how clearance-time dispersion becomes an appropriability loss, which is the step taken below.

Heterogeneity is what turns the queue into a source of variance. Device classes of very different complexity draw on the same limited pool of assessors, and the resulting service times are widely dispersed. Queueing theory relates waiting time to the squared coefficient of variation of service time, so a pool serving heterogeneous dossiers produces longer and more dispersed waits than a homogeneous system at the same traffic intensity (Kleinrock, 1975). Such dispersion is a feature of how technology is matched to procedure and does not disappear once a transitional backlog has been worked off.

Firms react to that dispersion as real-options theory predicts. Where submitting a dossier commits costs that cannot be recovered and the state of the queue becomes clearer over time, the option to defer entry gains value as the spread of clearance time widens (Dixit and Pindyck, 1994), which is an incentive to delay that exists independently of the mean. Entry-timing advantages in institutionally induced transitions accordingly differ between leaders and followers (Asimakopoulou and Paraskevopoulou, 2025). Antonelli (2017) describes an appropriability trade-off shaped by the architecture that governs how much generated knowledge a firm can capture; read against the MDR, the architecture of conformity assessment becomes a locus of appropriability policy in its own right.

A facially uniform architecture can nonetheless distribute its burden unevenly. The General Data Protection Regulation imposed identical requirements on all firms, yet its consequences for market concentration differed with firms' capacity to absorb compliance costs (Johnson et al., 2023). Proposition 2 derives the device counterpart, in which uniform per-dossier procedures meet a population that differs in size, documentation capacity and access to notified bodies.

One structural condition underlies the analysis. A regulatory regime translates between a technology that keeps

changing and the procedure used to validate it, and where the regime's translation properties do not match the technology, the mismatch propagates into firm value. For a conventional device it shows up along two margins. Development is largely deterministic while clearance is stochastic, which is why the variance of clearance time, and not only its mean, enters firm value (R1, with its sequencing consequence in P1). A population that is heterogeneous in firm size, device complexity and dossier completeness also meets uniform procedures, and heterogeneous input processed uniformly generates incidence that differs by firm type (P2).

Financing adds a third margin. Developers commit the cost of clinical evidence, quality-management systems (QMS) and technical documentation assessments (TDA) before certification, and smaller firms finance the wait from funds raised for a limited runway, so the probability that clearance outlasts the runway becomes the operative quantity for them. At the mean of 18.2 months implied by the twentieth survey, raising the coefficient of variation (CV) from 0.25 to 0.45 raises the probability that certification takes longer than 24 months from 10.5 to 19.4 percent, and US premarket approvals show how long such tails run in file-level data, with one decision in ten taking longer than 27 months (Table 1).

Patent-based appropriability defines the scope of the mechanism. Many Class IIb devices are incremental improvements that rely on know-how, speed and customer relationships more than on strong patents, and device generations often turn over before the patent term runs out, in which case the commercial window and not the statutory term binds (Holgersson, 2013; Somaya, 2012). The analysis therefore applies with full force to patent-protected standalone devices of Classes IIb and III, such as implants and novel therapeutic devices, and more weakly to incremental products.

Three expectations follow for the empirical analysis, and they answer the research question in turn. First, because a market-access procedure subtracts calendar time from a fixed term, the MDR should consume more appropriable term than the US routes and more than the directives it replaced, by an amount measured in years rather than months. Second, because the loss depends on the distribution and not only on its mean, a firm's decision on the order of market entry should turn on what completing the European file can save and on what a long European wait costs in financing, not on average durations. Third, because smaller and less experienced applicants have thinner regulatory departments and fewer completed files behind them, their clearance times should differ from those of large or frequent applicants more in the upper tail than in the mean, so that pooled statistics understate their exposure. Section 4 evaluates the first two with the empirical distributions and the two-market valuation and the third with the mixture identity, the tail-order result and an estimation on the file-level US record.

In the formal model the condition reduces to one measurable object, the variance $H_j \equiv \text{Var}[T_j]$ of effective patent term in jurisdiction j , which equals the variance of clearance time as long as expiry does not truncate the distribution; the welfare-relevant objects are tail functionals of the same distribution, because frictions and not variance as such make dispersion costly. The material that carries these quantities is assembled next.

3 DATA AND DESCRIPTIVE ANALYSIS

All inputs are publicly available, and the two sides of the comparison rest on records of different quality. European data come from the notified-body survey that the Commission has commissioned since 2023 (Austrian National Public Health Institute et al., 2026a,b), which continues a series the Commission itself collected from 2021 to October 2022, and from the manufacturer surveys of MedTech Europe (MedTech Europe, 2022, 2024, 2025); both report durations in brackets, the first by notified body and the second by firm, so every European bound reported below bounds the distribution of typical durations across bodies and not the distribution across dossiers. US data are file level. Two further asymmetries deserve notice: the clocks start at different points, the European interval at a signed written agreement and the US interval at receipt of the submission, and MedTech Europe reports that more than half of the European time falls outside the review phase itself. On the US side, the FDA publishes the date of receipt and the date of decision for every premarket notification and every premarket approval application, together with the applicant, the submission type, whether a third party reviewed it and the review advisory committee, so that the clearance time of each dossier can be computed directly and related to who submitted it and what it was (Food and Drug Administration, 2026a,b). Review-time goals agreed under the fifth Medical Device User Fee Amendments (MDUFA V) and the performance reports against them (US Food and Drug Administration, 2022, 2025, 2026) serve here only as context, because they count review time at the agency and exclude the periods in which a submission rests with the applicant. One measurement caveat follows from comparing the two US sources directly. Matching the Federal Register notices to the premarket approval database by approval date is unambiguous for 49 of the 70, since seven approval dates carry more than one original application and fourteen have no counterpart in the database. Among the 49 unambiguous links the submission date the notice states coincides with the database receipt date in 35 cases and falls 1 to 51 days later in 14, with a median gap of four days, and never earlier. Administrative duration as the database records it and the regulatory review period as the notice determines it are therefore closely related but not identical, which is one reason the clearance-time measurement and the restoration arithmetic are kept apart.

Survey rounds follow a fixed design. A small set of questions is asked every two months, a medium dataset covering activities since designation every four months, and a large dataset once or twice a year. All 53 bodies designated under the MDR or the in-vitro diagnostics regulation (IVDR) answered the twentieth round, 34 of them designated under the MDR only, 18 under both regulations and one under the IVDR only (Austrian National Public Health Institute et al., 2026b). Responses were complete in every round except the second and the third, for which values from earlier rounds were carried forward. Counts are cumulative from each body's designation and include applications for changes to existing certificates, so they measure the workload of the certification system and not the number of new devices.

3.1 Data sources and definitions of variables

Each survey round reports, cumulatively from designation onwards, how many MDR applications notified bodies have received and how many certificates they have issued, so the series traces a clearance system as it built capacity from a near-standstill in 2022. Online Resource 1, Table H.1 consolidates the record in a single vintage, the comparison published with the twentieth round. The application-to-certificate ratio stood near one quarter at the outset and reached 0.55 by February 2026; it is read below as a throughput signal and not as a measure of any individual device's waiting time.

Beyond throughput, the twentieth round documents where time accumulates. In 65 percent of files a written agreement followed the application within two months, so the dispersion does not sit mainly in contracting (Austrian National Public Health Institute et al., 2026b). From agreement to a new certificate, however, reported times spread widely. Among the 39 bodies that had issued combined quality-management and product certificates, 51 percent reported 13 to 18 months, 31 percent 19 to 24 months and 10 percent more than 24 months, whereas for quality-management certificates alone 62 percent reported 13 to 18 months and 30 percent 6 to 12 months. Combined certificates are the route that Class III devices and most implantable Class IIb devices must take, which places the widest dispersion exactly where the patent-term problem is most acute.

Incomplete documentation is part of the mechanism. Only 14 of the 52 MDR-designated bodies reported that more than half of the submissions they receive are complete on arrival, and of the 989 cumulative refusals recorded by February 2026 the two most frequent reasons were applications outside the scope of a body's designation (30 percent) and incomplete applications (25 percent) (Austrian National Public Health Institute et al., 2026b). Resubmission cycles lengthen the upper tail of the duration distribution and fall disproportionately on firms with thin regulatory departments (Huusko et al., 2023; Carl and Hochmann, 2024).

Two features of the record call for caution. Aggregates differ between datasets: the bimonthly small dataset reported 31,902 applications lodged by February 2026, against 32,898 in the annex sums of the medium dataset, and the source flags changes in counting methods at some bodies, for refusals in rounds 13 and 18 and for the annex sums of February 2026 (Austrian National Public Health Institute et al., 2026b). Reported durations also describe what a typical file takes at each body and not the distribution across files, so shares of bodies cannot be converted directly into probabilities for individual dossiers. The bracket method itself can be checked where file-level records exist. Fitted to six duration brackets of the 12,795 traditional 510(k) clearances of fiscal years 2021 to 2025, a log-normal recovers the mean closely (6.36 against 6.22 months) and overstates the coefficient of variation moderately (0.80 against 0.70), so bracket fits describe the distribution they are fitted to; they cannot turn body-level shares into file-level dispersion. Under the IVDR, 67 percent of the 12 bodies that reported typical durations for combined certificates gave 13 to 18 months and 8 percent gave 19 to 24 months, with none above 24 months, a narrower spread than under the MDR within the same institutional framework

(Austrian National Public Health Institute et al., 2026b). The comparison points to device-class composition and file characteristics as sources of dispersion and not to the framework as such, with the caveat that the IVDR population of bodies is small and its transition followed a different timetable.

Four further records complete the empirical base. The FDA publishes the receipt and decision dates of every 510(k), De Novo and premarket approval decision, from which the time from submission to decision follows file by file (Food and Drug Administration, 2026a,b). Federal Register notices that determine the regulatory review period for patent term restoration report for each device the dates of the investigational device exemption, of the premarket approval application and of approval, together with the patent concerned, whose filing date follows from the patent record (Federal Register, 2026). MedTech Europe’s 2024 survey reports certification fees and phase durations by firm size (MedTech Europe, 2024), and turnover by size class, cost of capital, valuation multiples and financing intervals come from Eurostat, Damodaran and Carta (Eurostat, 2026; Damodaran, 2026; Carta, 2025, 2026).

Comparing a low-risk European route with a high-risk American one would manufacture a gap out of classification. The author therefore pairs pathways that screen devices of similar clinical risk. Pairing by risk class is an assumption and not a matching procedure, since neither record reports the other jurisdiction’s classification for the same device, and the populations passing through the two routes differ in composition. Residual differences in device mix are absorbed into the reported gap, which is why the article reports the gap against two US comparators rather than one. The 510(k) route, which clears moderate-risk devices on a substantial-equivalence standard, serves as comparator for MDR Class IIb; premarket approval (PMA), reserved for the highest-risk devices, pairs with Class III. Holding clinical risk fixed isolates the procedural architecture as the source of any difference in clearance time. Table 1 reports what public records identify for each pathway, with European figures from the notified-body survey and firm-level survey averages and distribution-free bounds where only brackets exist.

MDUFA goals are negotiated commitments financed by user fees and reported against realised performance in annual reports to Congress, which is why the US side of the comparison can be anchored on public targets and outcomes (US Food and Drug Administration, 2022, 2026). Whether a comparable commitment would change behaviour is the kind of question that separates regulatory design from regulatory intent (Battaglion et al., 2023). No commitment of that kind existed on the European side during the observation period; the maximum timelines of Implementing Regulation (EU) 2026/977 apply only to procedures whose written agreement is signed on or after 25 February 2027 (European Commission, 2026).

Comparator completeness deserves explicit treatment. Specialised US pathways such as the Breakthrough Devices Program, De Novo classification and the humanitarian device exemption have no direct MDR equivalent beyond conditional certification of high-risk devices, which notified bodies rarely use (Dobrzynska et al., 2025; Gupte et al., 2025). Excluding those pathways from the comparison therefore probably understates the EU–US gap.




Fig. 2 Empirical clearance-time distributions. US curves are empirical distribution functions of the time from receipt to decision for substantially equivalent 510(k) decisions ($N = 15,706$) and original premarket approvals ($N = 311$) in the FDA databases. EU markers are cumulative shares of the 39 notified bodies by typical time to a combined quality-management and product certificate in the twentieth survey, with a log-normal fit. The dotted line marks the median interval of 696 days between financing rounds reported by Carta for the second quarter of 2025

Converting clearance-time distributions into the headline erosion and gap figures rests on three maintained assumptions. Clearance is taken as the event that binds effective term, which holds for standalone devices but not for combination products eligible for supplementary protection. Each device is assumed to rely on one core patent, so the device-level figure is an upper bound on what portfolio management cannot offset. Headline figures evaluate clearance at the conditional means, with variance entering separately through the mechanism formalised next. Qualitative findings survive changes in these assumptions; the precise headline numbers do not, and they are reported with that caveat.

3.2 Descriptive analysis of clearance times

Fig. 2 plots the empirical distribution functions of US clearance times against the cumulative shares of notified bodies by typical EU certification time. The US curves rise steeply and then flatten into long tails that the file-level records make fully visible, since one 510(k) clearance in ten takes longer than 10.4 months and one original premarket approval in ten longer than 27 months. For the EU, three reported points and a fit are all the public record offers. Before the MDR, European certification was commonly faster than US review; that advantage has turned into a US advantage for moderate-risk devices, and the EU has no counterpart to the US programme for breakthrough devices (Daigle and Torsekar, 2019; Gupte et al., 2025).

Cumulative throughput, the ratio of certificates to applications, rose from 0.25 in October 2022 to 0.34 in February 2024 and to 0.55 by February 2026 (Online Resource 1, Fig. H.1). Within the first five observations the marginal ratio of new certificates to new applications ranged between 0.36 and 0.54, consistent with uneven designation timing and lumpy additions of capacity; 39 bodies took part in the first study round (data to 31 March 2023) and 45 in the round for February 2024. Beyond February 2024 the marginal ratio stops measuring period throughput, because small application increments and revisions of earlier counts dominate its denominator (it exceeds one in two of the later intervals and turns negative in

Table 1 Clearance-time distributions by pathway from public records

Pathway and record	N	Mean	P50	P90	CV	Basis
EU MDR, typical time per body to a combined certificate	39 bodies	18.17	[13, 18]	> 24	0.26	Twentieth notified-body survey, shares 8/51/31/10 percent for 6–12, 13–18, 19–24 and more than 24 months (Austrian National Public Health Institute et al., 2026b); log-normal fit
EU MDR, firm-reported average, QMS / TDA	39–43 / 28	19.5 / 21.8	–	–	–	MedTech Europe (2024, p. 6)
EU MDD, implied	–	6.5–9	–	–	–	MDR certification of 13–18 months described as double the time under the directives (MedTech Europe, 2022)
FDA 510(k), substantially equivalent, FY2021–2025	15,706	5.51	4.37	10.38	0.80	FDA 510(k) database (Food and Drug Administration, 2026a)
FDA 510(k), traditional, no third party	12,795	6.22	5.16	11.27	0.70	Same
FDA De Novo, granted, FY2021–2025	173	12.12	10.35	21.45	0.57	Same
FDA PMA, original approvals, FY2016–2025	311	14.72	10.48	27.04	0.97	FDA PMA database (Food and Drug Administration, 2026b)

Note: Months. FDA durations run from receipt to decision in calendar time and include time on hold; decisions are assigned to fiscal years ending in September. QMS denotes the quality-management system audit, TDA the technical-documentation assessment and MDD the Medical Devices Directive that the MDR replaced. FDA quantiles use the inverse empirical distribution function throughout, as in Table 2. For the EU survey the brackets give distribution-free ranges; the median falls inside the closed bracket of 13 to 18 months, whereas the cumulative share of 0.897 at 24 months places the ninetieth percentile in the open upper bracket, which the data bound from below only. The log-normal fit gives a mean of 18.17 months, P50 of 17.61 and P90 of 24.28 months, and its CV of 0.26 describes typical durations across bodies, not files. The MDUFA V goals of 112 to 128 days for 510(k) and 285 to 290 days for PMA are themselves stated as average total time to decision in calendar days (US Food and Drug Administration, 2022). Reproduced by the script of Online Resource 1, Section F.

the last). What the diagnostic establishes is limited and deliberately so. The record is that of a capacity-bound system that expanded unevenly; the variance of individual clearance times cannot be read from cumulative counts and is reconstructed instead from the reported duration brackets.

Online Resource 1, Fig. H.3 summarises level and spread of clearance time by pathway. Measured by the mean, the MDR roughly doubled certification time relative to the MDD and placed EU certification far above the 510(k) route and above premarket approval. Two channels have to be kept apart when reading the figure. The mean channel moves the markers and produces the erosion measured in years, whereas the dispersion channel widens the bars and appears in the ninetieth percentile, in the probability of long waits and in the shortfall of Table 2, not in expected effective term. For the EU the bar describes the spread of typical durations across bodies, a lower bound on the spread across files, while the US bars describe files; the figure thus shows a clear gap in means and leaves the ranking of dispersion open.

Across pathways these quantities differ far more than means alone suggest. Taking the median interval of 22.9 months between financing rounds as the runway, the FDA records show exceedance probabilities of 0.7 percent for 510(k) clearances, 8.1 percent for De Novo requests and 13.8 percent for original premarket approvals, with expected shortfalls of 0.04, 0.55 and 2.41 months. For EU certification the survey brackets bound the exceedance probability between 10 and 41 percent and the shortfall between 0.11 and 1.66 months if the open bracket ends at 36 months (Table 2). Where dispersion enters through these channels it lowers value, and the pathway with the thinner tail relative to the runway is worth more for exactly the firms that cannot absorb a long wait.

Tail measures express the result in a form that a financing plan can use. One 510(k) clearance in ten takes longer than 10.4 months and one original premarket approval in ten longer than 27.0 months, and in the slowest tenth the average wait is 15.3 and 46.4 months. For EU certification the seventy-fifth percentile of typical durations lies between 19 and 24 months while the ninetieth lies above 24 with no upper bound in the data, and 41 percent of the bodies with combined certificates

reported typical durations above 18 months (Table 2). Among firms, half of the SMEs but a quarter of large manufacturers reported technical documentation reviews longer than 15 months (MedTech Europe, 2024).

4 ANALYSIS

Online Resource 1, Fig. H.2 summarises the causal chain that the formal apparatus captures. The architecture of conformity assessment generates a dispersed clearance clock, the dispersion translates into a distribution of effective patent term, frictions in risk bearing, financing and irreversible commitment convert that distribution into an appropriability loss, and differences in appropriability drive sequencing choices and, in the aggregate, an incidence that differs by firm type. Each box corresponds to a formal object defined below.

4.1 Estimation framework and identification

Consider the patent whose term is at issue, filed at date t_p , with statutory protection running $T^{\max} = 20$ years from that filing (Article 63(1) of the European Patent Convention (EPC); 35 U.S.C. §154(a)(2)); where a later application is the term-bearing one, the clock runs from its own filing date. The firm submits its dossier at $t_r \geq t_p$, and clearance in jurisdiction j requires a further random time τ_j with mean μ_j and standard deviation σ_j . Development before filing consumes no patent term; what counts is the calendar time between the term-bearing filing and market authorisation.

Definition 1 (Effective patent term). Let $D \equiv t_r - t_p \geq 0$ denote the interval between filing and regulatory submission. Effective patent term in jurisdiction j is

$$T_j = \max\{0, T^{\max} - D - \tau_j\}. \quad (1)$$

Two intervals are consumed before market entry, the deterministic preparation interval D and the stochastic clearance interval τ_j , and the maximum operator handles clearance after expiry. The innovator keeps whatever remains of twenty years once both have been subtracted, and because the second is random, so is the term that remains.

Table 2 Tail exposure of clearance time from public records

Pathway or group	Clearance time (months)		Probability of exceeding		Shortfall beyond	Mean of slowest
	P50	P90	18 months	22.9 months	22.9 months	tenth (months)
EU MDR, typical per body, from counts	[13, 18]	> 24	0.410	[0.103, 0.410]	[0.11, 1.66]	[24.0, 36.0]
EU MDR, typical per body, log-normal fit	17.61	24.28	0.465	0.149	0.49	27.49
EU MDR, TDA review phase, large firms, fit	12.6	19.5	0.15	0.04	—	—
EU MDR, TDA review phase, SMEs, fit	15.3	34.3	0.40	0.26	—	—
FDA 510(k), FY2021–2025	4.37	10.38	0.019	0.007	0.04	15.25
FDA De Novo, FY2021–2025	10.35	21.45	0.179	0.081	0.55	28.22
FDA PMA originals, FY2016–2025	10.48	27.04	0.228	0.138	2.41	46.36

Note: The financing interval of 22.9 months is the median of 696 days between funding rounds (Carta, 2025); shortfall is $E[(\tau - R)^+]$ from equation (3) and the last column expected shortfall at the 90 percent level. Quantiles and expected shortfall use the inverse empirical distribution function, so the last column is the integral of the step quantile function above the ninetieth percentile and not the mean of the observations above it, which for a discrete distribution is a different quantity. The printed shares of 8, 51, 31 and 10 percent among 39 bodies correspond to counts of 3, 20, 12 and 4, from which the share above 18 months follows exactly, because 18 is a bracket boundary, while the share above 22.9 months is bounded only within the bracket that contains it. The cumulative share at 24 months is 0.897, so the ninetieth percentile lies in the open upper bracket and the data give it no upper bound; the bracket in the last column assumes the open bracket ends at 36 months. Review-phase rows are log-normal fits to the firm-size shares of MedTech Europe (2024, Fig. 8).

A worked example fixes ideas without requiring the interval D to be known. Certification at the mean of the survey-based EU distribution consumes 1.51 years of term, and at its ninetieth percentile at least two years, since the survey counts place that percentile above 24 months. An original premarket approval consumes 1.23 years at the mean and 2.25 at the ninetieth percentile, a reminder that the US tail is long as well; a 510(k) clearance consumes 0.46 and 0.87 years (Table 2). Differences between routes are what the records identify. Levels of residual term would additionally require the date from which the term-bearing patent runs, and for devices whose patents descend from earlier applications that date is the earliest application in the chain of benefit rather than the filing date of the granted patent; establishing it for a sample of devices is a task in patent-register work that this article does not undertake, so no absolute effective term is reported.

Throughout, D is treated as predetermined, although a firm can trade calendar time before submission against the risk of an incomplete file. Preparing longer consumes term deterministically, whereas submitting early risks resubmission rounds that lengthen the right tail of τ . The erosion figures therefore isolate the clock after submission and omit a privately chosen consumption of term that is more attractive for well-financed applicants and less feasible for SMEs whose runway is short. Proposition 1 captures one version of that trade-off, and Online Resource 1, Section B.5 the general case.

Definition 2 (Clearance-time variance). For jurisdiction j , the variance of effective patent term at clearance is $H_j \equiv \text{Var}[T_j]$. As long as expiry does not truncate the distribution, $H_j = \sigma_j^2$. The EU–US gap in clearance-time variance is $\Delta H_{EU/US} \equiv \sigma_{EU}^2 - \sigma_{US}^2$.

Clearance-time variance is a structural property of the regime and not a state of the innovator’s knowledge: a better-informed firm could resolve epistemic uncertainty about what a regulator will decide, whereas the dispersion considered here persists in steady state and is priced even by a firm with correct expectations. Two systems can share an identical μ and differ sharply in H .

Queueing enters only as an architectural argument. For a single server with Poisson arrivals and general service times, the Pollaczek–Khinchine relation gives the expected sojourn

time

$$W = E[S] + \frac{\rho E[S] (1 + CV_S^2)}{2(1 - \rho)}, \quad (2)$$

with $E[S]$ the mean service time, CV_S its coefficient of variation and ρ the traffic intensity (Kleinrock, 1975). Only the comparative statics are used, since at a given traffic intensity a pool serving heterogeneous device classes generates longer waits than one with standardised processing. Levels of W are not used, because a single server is a stylised stand-in for the 52 designated bodies.

Three sources of dispersion can be separated with standard results. A mix of device classes in one pool adds a between-class component to the variance of service times, by the same law of total variance that underlies Proposition 2, and the consultation step of Article 54 MDR adds an externally paced stage for part of Class III. Resubmission compounds the clock. If each round has coefficient of variation CV_S and is followed by another with probability r , the squared coefficient of variation of total clearance time is $(1 - r)CV_S^2 + r$, so a per-round CV of 0.25 becomes 0.33 to 0.40 at resubmission probabilities of 5 to 10 percent, and matching the file-level CV of 0.70 measured for traditional 510(k) clearances would require about 46 percent (Online Resource 1, Section D). Congestion adds a transitional component on top, which the survey-based service distribution puts at a coefficient of variation of 0.84 at a traffic intensity of 0.85 and 0.44 at 0.30, against the service-time value of 0.26, so the part of today’s dispersion that stems from class mix, consultation and resubmission does not disappear once the backlog clears. Dispersion between bodies is visible directly in the twentieth survey, where the typical durations reported by 39 bodies imply a standard deviation across bodies of between 2.6 and 8.2 months if the open bracket ends at 36 months.

A firm’s filing and sequencing decision is a stochastic control problem. Residual term decays deterministically, clearance arrives as a random event whose hazard depends on how long the dossier has been pending, and the firm chooses among filing strategies under a finite horizon. The natural formal object is a piecewise-deterministic Markov process whose value function satisfies a Hamilton–Jacobi–Bellman equation with jump terms (Davis, 1993; Fleming and Soner, 2006); the state space, the age-dependent hazard, the equation and the sequencing condition are set out in Online Resource 1, Section B.

How dispersion enters value has to be stated with care. For a risk-neutral firm discounting a constant revenue flow from clearance until expiry at rate δ , present value at submission per unit of flow is $(e^{-\delta\tau} - e^{-\delta T^{\text{res}}})/\delta$, convex in τ , so dispersion does not lower expected value; at the survey-based EU parameters and the European healthcare cost of capital of 9.57 percent it raises it by the equivalent of about 2.6 days of shorter delay (Damodaran, 2026). A diversified investor without financing constraints would barely care about the spread.

Dispersion becomes costly through four channels that the benchmark omits. Owners who cannot diversify a single-product firm bear the variance of project value as a risk premium. Financing that has to bridge the waiting time fails when clearance outlasts the runway. Commitments made before clearance, such as production capacity, clinical staff and distribution agreements, lose value when their timing cannot be adjusted, which raises the value of waiting (Dixit and Pindyck, 1994). Commercial windows fixed by product life cycles or by competitor entry cap the period over which a delay can be recovered. For a financing runway of R months, the probability of running out before certification and the expected shortfall are

$$P(\tau > R), \quad E[(\tau - R)^+] = E[\tau]\Phi(d_1) - R\Phi(d_2), \quad (3)$$

where, for a log-normal clearance time with parameters $(\mu_{\ln}, \sigma_{\ln})$, $d_2 = (\mu_{\ln} - \ln R)/\sigma_{\ln}$ and $d_1 = d_2 + \sigma_{\ln}$.

Every parameter the author uses comes from a public record, and Table 3 lists them with their sources. Clearance-time distributions follow Table 1, with the log-normal fit to the twentieth survey as the European central case and the firm-reported average of 21.8 months as the upper case. The valuation horizon T^{res} is set at eleven years, a plausible residual term for a device entering review, and it is not an estimate: the sequencing thresholds below change by at most 0.11 months when it is varied between eight and twenty years, because expiry truncation almost never binds at those horizons. Three inputs carry more uncertainty than their decimals suggest. The European distribution is fitted to four bracket shares reported by 39 bodies, so its coefficient of variation of 0.26 is an upper bound on precision rather than a measurement, since the same procedure applied to file-level 510(k) data overstates dispersion by about sixteen percent. Turnover per enterprise stands in for the annual revenue of a single device and overstates it for multi-product firms, so the cost ratios are lower bounds on the burden per device and upper bounds on the burden per firm. And the mixture weights come from the SME share of certificates while the group moments come from shares of survey respondents, two universes that coincide only if response rates do not vary with firm size. Because the interval between the term-bearing filing and submission is not established here, the article reports differences between pathways throughout, which do not depend on it. Online Resource 1, Table H.4 varies each input across its documented range.

Where the record identifies the applicant, the third expectation can be estimated rather than described. For each substantially equivalent premarket notification i decided in fiscal years 2021 to 2025, let $g(i)$ assign the applicant to one of four submission-frequency classes by the number of such notifications the same applicant obtained in that window, a single

submission, two to four, five to nineteen, or twenty and more. The estimating equation is

$$\ln \tau_i = \alpha + \sum_{g \neq G} \beta_g \mathbf{1}[g(i) = g] + \gamma' x_i + \delta_{t(i)} + \theta_{a(i)} + \varepsilon_i, \quad (4)$$

where the reference class G is twenty and more submissions, x_i contains the submission type, an indicator for third-party review and an indicator for a US applicant, δ_t are fiscal-year effects and θ_a fixed effects for the twenty review advisory committees that stand in for device area, with standard errors clustered by applicant. Because the argument concerns the tail rather than the centre, the same specification is estimated as a quantile regression at the quartiles and at the ninth decile, so that $\beta_g(q)$ traces how the experience gap changes along the distribution. Submission frequency is a descriptive classifier rather than a measure of experience at the moment of filing: it counts decisions across the whole window, including decisions that postdate the submission at hand, so a long review can itself affect how many completed procedures a firm shows. It is also not a measure of firm size. The coefficients are therefore read as conditional associations between clearance time and how often a firm appears in the register, not as effects of accumulated experience. The first expectation implies $\beta_g > 0$ for every class below the reference, the third implies $\beta_g(0.9) \geq \beta_g(0.5)$, and as a distribution-free complement the tail-order condition of Proposition 3 is checked by comparing the empirical distribution functions of the least and the most experienced classes at every point of a fine grid.

4.2 Results

Four findings follow, and they claim different things. A descriptive comparison of effective term across pathways comes first (R1), followed by a sequencing condition that prices dispersion through dossier completion and financing frictions (P1), an estimation of the experience gap on the file-level US record, and a mixture-variance identity that locates a disproportionate share of dispersion with SMEs (P2).

Measured against the directives, the MDR consumes additional patent term equal to the lengthening of certification. MedTech Europe reported MDR certification at 13 to 18 months on average and described this as double the time historically needed under the directives (MedTech Europe, 2022), which implies 6.5 to 9 months of additional term consumed, or 0.54 to 0.75 years. Against the US comparators the gap depends on the pathway. With EU certification at 18.2 months (survey fit) to 21.8 months (firm-reported technical documentation assessments) and 510(k) clearance at 5.51 months, the transatlantic difference is 12.7 to 16.3 months, or 1.05 to 1.36 years; against original premarket approvals at 14.72 months it is 3.5 to 7.1 months, and against their median of 10.48 months 7.7 to 11.3 months. Whether the EU clock is also more dispersed is a separate question that public data do not settle. The FDA records give file-level standard deviations of 4.38 months for 510(k) clearances and 14.30 months for premarket approvals. For the EU only the dispersion of typical durations across bodies is available, between 2.6 and 8.2 months distribution-free, and firm-reported review phases with standard deviations of 4.6 months for large manufacturers and 13.1 months for SMEs under a log-normal fit, and neither measures the spread across files.

Table 3 Valuation and financing inputs from public records

Parameter	Value	Source
EU submission cost K	Notified-body fees of 136,981 euros for a QMS and 176,202 euros for a TDA certificate, 313,183 euros together; total first-year certification cost of about 4.47 million euros, since fees make up 7 percent of it	MedTech Europe (2024, p. 6)
Annual revenue π , proxy	Net turnover per enterprise, manufacture of medical and dental instruments and supplies, class C32.5 of the European classification of economic activities, EU27, 2023: 1.14 million euros (10–19 persons employed), 3.21 million (20–49), 16.58 million (50–249)	Eurostat (2026)
Cost ratio K/π	0.02–0.28 years with fees, 0.27–3.94 years with total cost	Derived
Cost of capital δ	7.54 percent (US healthcare products, weighted average cost of capital, WACC), 9.57 percent (European healthcare, WACC), 10.57 percent (European healthcare, cost of equity)	Damodaran (2026)
Financing interval R	696 days (22.9 months), median time between new funding rounds, second quarter of 2025	Carta (2025)
Cost of a further round c	Median dilution of 16 percent in 2025 and 19 percent two years earlier, times enterprise value to sales of 3.45 (European healthcare) or 4.76 (US healthcare products): 0.55 and 0.90 years of revenue	Carta (2026); Damodaran (2026)
Regulatory review period	Testing phase median 5.01 years (60 notices); approval phase median 0.86 years (70 notices)	Federal Register (2026)
SME share of certificates	40 percent	MedTech Europe (2022)
TDA review phase by firm size	More than 15 months: SMEs 50 percent, large manufacturers 25 percent; 9–15 months: 30 and 56 percent; 4–8 months: 20 and 11 percent	MedTech Europe (2024, Fig. 8)

Note: Turnover per enterprise includes exports and all products, so K/π overstates the burden for multi-product firms and understates it where a device earns only part of turnover. All inputs and derivations are part of the reproduction package (Online Resource 2, documented in Online Resource 1, Section F).

In other words, the European route costs a moderate-risk device about a year of exclusivity on average relative to a 510(k) clearance, while a Class III device loses less against a premarket approval, whose own tail is long. What the European record cannot show is how uncertain the surviving exclusivity is for an individual file; that gap in the evidence, and not a measured variance ratio, is the empirical starting point for the reporting argument below. Both findings are consistent with evidence that finite and uncertain terms depress investment in long-development innovation (Czarnitzki and Toole, 2011; Budish et al., 2015).

A distributional caveat applies. The European figures derive from a log-normal fit to four bracket shares whose last bracket is open above 24 months, so heavier-tailed families cannot be distinguished from the log-normal with these data. On the European brackets the log-normal, gamma and Weibull families agree on the coefficient of variation to within three percent, because the reported shares concentrate in two adjacent brackets; on the file-level 510(k) brackets, where the truth is known, the log-normal overstates dispersion by 16 percent and the other two families understate it slightly (Online Resource 1, Section A). Any misspecification therefore bears on the European tail, which the article bounds rather than fits, and not on the ordering of pathways by mean.

Whether a developer should enter the EU and the United States at the same time or the United States first depends on what the firm can do while the US review runs. Consider a firm whose European sales yield a constant flow π from clearance until expiry, whose EU submission costs a lump sum K , and whose funds last for a financing interval R , so that each further interval spanned by the European wait requires a financing round that costs the existing owners $c\pi$. Ignoring expiry truncation, the value of an EU submission at the moment it is made, when clearance time follows the distribution F , is

$$V(F) = \frac{\pi}{\delta} E_F[e^{-\delta\tau}] - K - c\pi E_F[\lfloor \tau/R \rfloor]. \quad (5)$$

Parallel entry submits the European file as it stands, with clearance distribution F_0 . US-first entry submits in the EU after US clearance, at time τ_{US} , and can use the US review to

complete the European file, which then clears according to F_1 .

Proposition 1 (Sequencing condition). *Let US and EU clearance times be independent and value in the United States be the same under both strategies. US-first entry is preferred to parallel entry if and only if*

$$E[e^{-\delta\tau_{US}}] V(F_1) > V(F_0), \quad (6)$$

and to entry in the United States alone if and only if $V(F_1) > 0$. If the firm would submit the same European file in either case ($F_1 = F_0$), US-first entry is never strictly optimal: parallel entry dominates it when $V(F_0) > 0$, and entry in the United States alone does when $V(F_0) < 0$.

Online Resource 1, Section B.5 gives the derivation; expiry truncation, which makes a later submission worth less, strengthens the last statement. Deferring a profitable European submission only delays its revenue, and deferring an unprofitable one is worse than not making it, so waiting has no value of its own for a firm whose European file would not change. Sequencing pays when the months of the US review are used to complete the European file, so that it clears faster and with a thinner tail, or when financing costs make the European wait expensive; dispersion enters through both routes.

Table 4 evaluates the condition with the empirical inputs of Table 3. With an unchanged European file, parallel entry beats US-first entry by 0.38 years of European revenue behind a 510(k) clearance and by 0.94 years behind a premarket approval. A completed file changes the balance only if it shortens European certification enough. Behind a 510(k) review, the required reduction of the mean is 5.3 to 5.5 months when only notified-body fees count and no refinancing is needed; it falls to 4.3 to 4.4 months when each financing interval spanned costs a round at the 2025 median dilution and European valuation multiples, to 3.7 to 3.9 months at the dilution of two years earlier and US multiples, and to 2.0 months for a firm of 10 to 19 persons bearing the full first-year certification cost. Behind a premarket approval the requirement rises to between 6.2 and 14.1 months, because

Table 4 Reduction of mean European certification time required for US-first entry (months)

		Cost of a further financing round c (years of revenue)		
US pathway and cost basis	Persons employed	0	0.55	0.90
<i>US route: 510(k) clearance</i>				
Notified-body fees	10–19	5.31	4.28	3.72
Notified-body fees	20–49	5.41	4.38	3.81
Notified-body fees	50–249	5.46	4.42	3.85
Total first-year cost	10–19	3.12	2.36	2.01
Total first-year cost	20–49	4.64	3.68	3.17
<i>US route: premarket approval</i>				
Notified-body fees	10–19	13.70	12.55	11.80
Notified-body fees	20–49	13.96	12.81	12.07
Notified-body fees	50–249	14.08	12.93	12.19
Total first-year cost	10–19	8.16	6.96	6.21
Total first-year cost	20–49	12.03	10.86	10.11

Note: Entries are the reduction of the mean of the European clearance distribution, at constant CV, that equates the values of US-first and parallel entry; larger reductions favour US-first entry. Inputs from Table 3: European healthcare cost of capital of 9.57 percent, financing interval of 22.9 months, K/π from notified-body fees or total first-year certification cost relative to turnover per enterprise; valuation horizon eleven years with expiry truncation, to which the entries are insensitive. European clearance follows the log-normal fit to the twentieth survey and US clearance times are resampled from the file-level FDA records; 300,000 draws. With an unchanged European file, US-first entry falls short of parallel entry by 0.38 years of EU revenue behind a 510(k) review and by 0.94 years behind a premarket approval. Results for cost of capital of 7.54 and 10.57 percent are given in Online Resource 1, Section B.5.




Fig. 3 When entering the United States first pays. Curves give the reduction of mean European clearance time that would make US-first entry exactly as good as parallel entry, at the European healthcare cost of capital; bands span the cost bases of Table 3, from notified-body fees alone to the full first-year certification cost, and the three firm-size classes within each. The dashed line marks the difference between the fitted mean review phases of small and large manufacturers. It is a cross-sectional gap between firm types and not a measured effect of preparing a dossier; it serves only as a quantity of comparable order against which the threshold can be read. Behind a 510(k) review the requirement meets that line without refinancing and falls below it with refinancing; behind a premarket approval it stays far above

that review itself takes 14.7 months on average. No public record measures what completing a European dossier during a US review actually saves, so the threshold has to be read against quantities of a comparable order rather than against an estimated effect. One such quantity is the gap of 5.3 months between the fitted mean review phases of small and large manufacturers (MedTech Europe, 2024); another is the 77 percent of respondents who report faster assessment for a second technical-documentation certificate than for the first. Both are of the order the 510(k) threshold requires and neither approaches what a premarket approval requires, which is the asymmetry the article reports; neither is evidence that file completion causes a saving of that size.

Fig. 3 reads the same condition off a single panel. The

order of entry therefore turns on two questions a project team can answer from its own files and from public data, namely how many months European certification would lose if the file were completed while a 510(k) review runs, and how many financing rounds a long European wait would require. Across the documented range of the cost of capital, 7.5 to 10.6 percent, the requirement moves by at most 0.9 months behind a 510(k) review and by up to 2.3 months behind a premarket approval, and the tail of the European clock, which a mean-based rule cannot see, enters through the refinancing term.

As a model-implied condition, Proposition 1 is not empirically identified; testing it would require matched records pairing patent filing dates with FDA submissions and MDR applications, which are not in the public domain. Industry surveys describing a shift of launch priority away from the EU are compatible with the mechanism (MedTech Europe, 2025; Svempe, 2024), yet other factors could contribute jointly, among them the pandemic shock to European manufacturing SMEs (Juergensen et al., 2020), the pricing of cross-Atlantic risk and heterogeneity in discount rates. Informational spillovers from US clearance are not modelled; they could shorten the European clock where clinical data from a premarket approval can be reused, and they are less plausible for 510(k) clearances that rest on substantial equivalence. Online Resource 1, Section E translates the condition into a decision tree for the project-valuation stage.

Proposition 2 (SME amplification). *Let large firms face clearance-time parameters (μ_L, σ_L) and SMEs (μ_S, σ_S) , with $\mu_S > \mu_L$ and $\sigma_S > \sigma_L$, and let a share p of applications originate from large firms. Portfolio-level variance is*

$$H^{\text{port}} = p\sigma_L^2 + (1-p)\sigma_S^2 + p(1-p)(\mu_L - \mu_S)^2. \quad (7)$$

Its last term is a between-group amplification that is strictly positive whenever $\mu_L \neq \mu_S$. It is a structural identity of the mixture and not a behavioural response, since it follows from the law of total variance applied to two groups with different means, and it holds for any group-level distributions with finite first two moments. Notified-body and industry statistics pool across firm types and therefore measure H^{port} , not

the dispersion each group experiences. The between-group term does not change under proportional reductions of within-group dispersion and depends only on the mean gap, so capacity expansion that raises aggregate throughput without closing the gap between firm types leaves it untouched (proof and examples in Online Resource 1, Section C). Proposition 3 below is accordingly stated for levels strictly above the crossing point, since at the crossing point itself an atom can reverse the order.

Before the European firm-size data are read through the identity, the US record allows the firm-level pattern to be estimated. Table 5 reports equation (4). Without controls, applicants with a single premarket notification in the window take 66 percent longer than applicants with twenty or more; with submission type, third-party review, applicant country, fiscal year and device area held fixed, the gap is still 36 percent, and it declines monotonically with experience, to 18 percent for two to four and 14 percent for five to nineteen submissions. Quantile regressions show where the gap sits. For single-submission applicants the coefficient rises from 0.196 log points at the first quartile to 0.295 at the median and 0.360 at the third quartile before settling at 0.348 at the ninth decile, so the least experienced applicants lose more at the upper end of the distribution than at its centre, which is the third expectation. In months, the ninth decile of clearance time is 13.4 for single-submission applicants against 8.6 for the most frequent, the mean of the slowest tenth 18.6 against 11.9, and the share exceeding eighteen months 4.1 against 0.7 percent (Fig. 4). The empirical distribution function of single-submission applicants lies below that of frequent applicants everywhere above about two months, so the tail-order condition of Proposition 3 holds on this record over the whole region of interest; at a handful of early jump points the order reverses by at most 0.004 in probability, which is why the condition is stated for levels above the crossing point rather than for the whole support. A Kolmogorov–Smirnov statistic of 0.266 rejects equality of the two distributions at any conventional level. Restricting the sample to traditional submissions without third-party review, dropping the backlog years 2022 and 2023, or moving the reference class to ten and more submissions leaves the single-submission coefficient between 0.245 and 0.306. Submission frequency is a proxy, and the association is conditional rather than causal; what the estimate establishes is that the pattern the European survey shows for SMEs appears in a file-level record with the same sign, the same monotonicity and the same concentration in the tail.

Two extensions answer the objection that the classification counts decisions the applicant had not yet obtained when the focal file was submitted. The first replaces it with a strictly backward-looking measure, namely the number of substantially equivalent decisions the same applicant had already obtained, over the whole public record back to 1996, strictly before that submission was received. Applicants entering with no prior decision wait 22.0 percent longer than those with twenty or more, those with one to four 13.7 percent longer and those with five to nineteen 8.4 percent longer, with the same controls and clustering (Online Resource 1, Table D.1); at the median the gap for a first-time applicant is 0.192 log points and at the ninth decile 0.178, so the pattern

is not an artefact of counting later decisions. The second holds the firm constant, and it changes the reading. Among applicants with at least three submissions, adding applicant fixed effects and regressing log duration on the log of the running submission number gives -0.433 (0.021) when the sample is defined by decision year. Right-censoring produces that estimate, since files received late and still open are absent from a decision-year sample, so late submissions are selected towards short durations and the bias runs with the finding rather than against it. Defining the sample by receipt cohort attenuates that selection, and the gradient collapses by about nine tenths, to -0.045 (0.021) for the 7,518 files received in fiscal years 2018 to 2021 and -0.016 (0.022) for the 7,456 received in 2019 to 2022, against a placebo on the reversed submission order of -0.011 (0.024). Only the first of the three differs from zero at conventional levels, its interval running from -0.086 to -0.003 , and a doubling of the submission count shortens clearance by about three percent against twenty-six under decision cohorts. Two limits attach to the exercise. Conditioning on the decision year instead of the receipt year within the same cohort restores a coefficient of -0.390 (0.021), which is what conditioning on an outcome does, since the decision year is the receipt date plus the duration being explained. And a register of decided files is not a complete receipt cohort, because applications still open or withdrawn never enter it. Within a firm, later submissions are therefore at most marginally faster, by an amount the placebo cannot separate from noise. The cross-sectional gap between frequent and infrequent applicants separates firms rather than tracking experience accumulating inside one, which is the narrower reading the article adopts.

Two objections arrive before the European evidence can be used at all, and both can be answered with the records in hand rather than deferred to future data. The first is that the survey describes notified bodies and not files, so that nothing follows about the distribution an individual applicant faces. Right as a statement about identification, the objection is wrong as a statement about direction. Writing τ for the duration of a file and b for the body that handles it, the law of total variance gives $\text{Var}[\tau] = E[\text{Var}[\tau | b]] + \text{Var}[E[\tau | b]]$, so the dispersion of body-level durations is a lower bound on the dispersion of file-level durations whenever the reported typical duration tracks the central tendency of a body's own distribution. Over every assignment of typical durations consistent with the reported counts, the standard deviation across bodies lies between 2.57 and 8.23 months if the open bracket ends at 36 months. Turning that into a bound on file-level dispersion needs one further assumption, since the survey counts bodies while a file distribution weights them by caseload: were caseloads equal, the between-body term would be at least 2.57 months and the file-level coefficient of variation at least 0.14 at the survey mean. Caseloads are unequal and unpublished, so the figure is conditional on that weighting rather than a bound the data deliver on their own. The second objection is that the two European sources measure different things. They do, and the difference is of the size the sources themselves imply: bodies report 18.17 months from a signed agreement, firms 21.80 months from application, a gap of 3.63 months that falls inside the pre-review phase which MedTech Europe puts at three to six months. The records are therefore not in

Table 5 Clearance time and applicant experience: substantially equivalent 510(k) decisions, fiscal years 2021–2025

Dependent variable: log clearance time	OLS	Quantile regression			
	mean	q=0.25	q=0.50	q=0.75	q=0.90
Single submission	0.306 (0.025)	0.196 (0.018)	0.295 (0.021)	0.360 (0.023)	0.348 (0.025)
Two to four submissions	0.163 (0.025)	0.104 (0.017)	0.135 (0.020)	0.236 (0.021)	0.187 (0.024)
Five to nineteen submissions	0.133 (0.025)	0.089 (0.016)	0.116 (0.019)	0.202 (0.021)	0.154 (0.023)
Twenty or more submissions		reference			
Special 510(k)	-1.248 (0.055)				
Third-party review	-1.835 (0.105)				
US applicant	-0.045 (0.015)				
Fiscal-year and device-area effects	yes	yes	yes	yes	yes
Observations	15,702	15,702	15,702	15,702	15,702
Applicants (clusters)	6,146				
R ²	0.418				
<i>Robustness of the single-submission coefficient, OLS</i>					
Traditional submissions, no third party	0.306 (0.027)	N = 12,793			
Excluding fiscal years 2022 and 2023	0.298 (0.028)	N = 9,290			
Reference class ten or more submissions	0.245 (0.021)				

Note: OLS denotes ordinary least squares. Coefficients on applicant experience classes with standard errors in parentheses, clustered by applicant for the least-squares column; quantile-regression standard errors from the asymptotic covariance. Clearance time is the calendar time from receipt to decision. Submission type includes traditional, special and abbreviated 510(k)s; device area is the review advisory committee (twenty categories). Applicant names are normalised before counting. Reproduced by the script of Online Resource 1, Section F.




Fig. 4 Clearance time by applicant experience. Empirical distribution functions of the time from receipt to decision for substantially equivalent 510(k) decisions of fiscal years 2021 to 2025, by the number of such decisions the applicant obtained in that window. The curve for single-submission applicants lies below every other curve at every point, so their clearance times dominate in the upper tail; the gap between the least and the most experienced class is 2.6 months at the median and 4.8 months at the ninth decile

conflict; they bracket the same procedure from two ends.

Pooled statistics understate SME-specific dispersion exactly when $\sigma_S^2 > \sigma_L^2 + (1-p)(\mu_L - \mu_S)^2$. MedTech Europe’s firm-size data allow the identity to be evaluated for the review phase of technical documentation assessments, the phase in which files are examined and deficiencies raised (MedTech Europe, 2024, Fig. 8). One feature of the source has to be recorded before the fits are used. The four shares reported for SMEs sum to 100 percent, those for large manufacturers to 92, and for the quality-management phase to 82 and 97; the missing mass is unexplained in the source and is here distributed proportionally. Log-normal fits to the shares so normalised give a mean of 13.4 months with a standard deviation of 4.6 months for large manufacturers and 18.7 months with 13.1 months for SMEs, and SMEs account for 40 percent of the certificates required (MedTech Europe, 2022), so that $p = 0.6$. Assigning the missing eight points of large-manufacturer mass entirely to one bracket instead moves the

between-group share between 4.7 and 8.8 percent and the pooled ninetieth percentile between 25.0 and 29.0 months (Online Resource 1, Section C), so the qualitative conclusion survives the ambiguity while the point estimates do not carry three significant digits. The condition holds by a wide margin (170.6 against 32.7 squared months), yet the between-group term contributes only 7.6 percent of pooled variance (Online Resource 1, Fig. H.4). Two caveats temper the numbers. The survey samples are small, with 28 respondents for technical documentation assessments, and total certification times hardly differ by firm size, because SMEs report shorter pre-review phases (MedTech Europe, 2024). A longer average wait is therefore not the main form of the SME burden in these data; that burden sits in the within-group tail of the review phase and in costs that are fixed per certificate. Once notified bodies report processing times cross-tabulated by firm size, the identity can be taken to administrative data.

Tail exposure differs more sharply than means. Half of the SME respondents but a quarter of large manufacturers reported review phases longer than 15 months, and under the log-normal fits the ninetieth percentile of the review phase is 34.3 months for SMEs against 19.5 months for large firms (Table 2). Fixed costs compound the difference. Notified-body fees of 313,183 euros for a quality-management and a technical-documentation certificate equal 0.28 years of turnover for the average enterprise with 10 to 19 persons employed in the manufacture of medical and dental instruments and supplies, 0.10 years with 20 to 49 persons and 0.02 years with 50 to 249 persons (MedTech Europe, 2024; Eurostat, 2026). Fixed compliance costs of that shape are what the literature on administrative burden identifies as the channel through which procedure weighs on smaller firms (Tomasi et al., 2023). Pooled statistics average both the tail and the cost ratio away, which is the practical meaning of the within-group and between-group terms taken together.

Three statements need to be kept apart, because they call for different policies. A longer mean clearance time for SMEs plausibly reflects access to bodies, file completeness and pre-submission contact, and responds to measures that improve

first submissions. Wider within-group spread plausibly reflects thin regulatory departments and mismatches between devices and designation scopes. The between-group term is an accounting consequence of pooling, so reporting by firm size makes it visible without changing either of the first two.

Online Resource 1, Table H.4 collects the documented alternatives behind the main quantities, and Section D of the same resource the alternative distributional families and service-time assumptions. None reverses the ordering of pathways by mean clearance time, none changes the conclusion that US-first entry behind a 510(k) review requires an improvement of the European file of the size documented between firm types, and varying the interval between filing and submission shifts all effective terms by the same amount.

Once dispersion is costly mainly through the tail, a tail measure of effective term summarises a clearance regime better than its variance. For a confidence level α and clearance time measured in years, effective term at risk is the protection that a dossier retains at the α -quantile of clearance time, and expected term shortfall is the protection retained on average by the slowest share $1 - \alpha$ of dossiers,

$$\begin{aligned} \text{ETaR}_\alpha &= T^{\max} - D - q_\alpha(\tau), \\ \text{ETS}_\alpha &= T^{\max} - D - \text{ES}_\alpha(\tau), \end{aligned} \quad (8)$$

where $\text{ES}_\alpha(\tau) = (1 - \alpha)^{-1} \int_\alpha^1 q_u(\tau) du$ is the expected shortfall of clearance time. The construction borrows from financial risk measurement, in which expected shortfall is the coherent counterpart of value at risk (Artzner et al., 1999; Acerbi and Tasche, 2002); to the author's knowledge, neither measure has so far been applied to regulatory clearance time or to patent term.

Proposition 3 (Tail concealment under pooling). *Let F_L and F_S denote the clearance-time distributions of large firms and SMEs, and $F = pF_L + (1 - p)F_S$ the pooled distribution with $0 < p < 1$. Suppose that SME clearance times dominate in the upper tail, $F_S(x) \leq F_L(x)$ for all $x \geq x_0$. Then for every $\alpha \geq \alpha_0 \equiv \max\{F_L(x_0), F_S(x_0)\}$*

$$\begin{aligned} q_\alpha^L &\leq q_\alpha \leq q_\alpha^S, \\ \text{ES}_\alpha^L &\leq \text{ES}_\alpha \leq \text{ES}_\alpha^S, \end{aligned} \quad (9)$$

so that pooled effective term at risk and pooled expected term shortfall overstate the protection that SMEs retain in the tail.

A proof of a few lines is given in Online Resource 1, Section G. Unlike Proposition 2, the result needs neither moments nor a distributional family; it places the pooled tail between the group tails at every level above α_0 . For the review phase of technical documentation assessments, the reported shares satisfy the condition above 15 months, where 50 percent of SMEs against 25 percent of large manufacturers are still waiting, but not at 8 months, by which a larger share of SMEs than of large manufacturers has finished; the log-normal fits cross at 10.2 months, so that α_0 is about 0.26 and the ordering covers the upper three quarters of the distribution. A reader of notified-body statistics can draw the consequence directly: a figure describing how long the slowest tenth of all certifications takes will always make the slowest tenth of SME certifications look faster than it is, whatever the shape of the

distributions, as long as SME files are the slower ones in the tail.

In these data the differences are material. The pooled 90th percentile of the review phase lies at 25.0 months, against 34.3 months for SMEs and 19.5 months for large firms, and at the 75th percentile the pooled value of 17.9 months compares with 23.4 and 15.9 months. Pooled statistics therefore credit SMEs in their slowest decile with more than nine months of protection that they do not retain (Online Resource 1, Section G, Table G.2).

Reported brackets identify these measures only partially (Manski, 2003). Without any distributional assumption, the α -quantile lies in the bracket in which the cumulative share crosses α , and expected shortfall is bounded by placing the tail mass at the lower or upper edges of its brackets (Online Resource 1, Section G). For typical EU certification times the median lies between 13 and 18 months and the 75th percentile between 19 and 24, but the ninetieth percentile is not bounded above at all: with counts of 3, 20, 12 and 4 among 39 bodies the cumulative share at 24 months is 0.897, so the quantile falls in the open bracket. Expected shortfall inherits that failure. Its lower bound is 24 months, whereas the upper bound equals whatever cap is assumed, 36 months at a cap of 36 and 60 at a cap of 60, and the log-normal value of 27.5 months lies inside that region without being implied by it. The FDA records need no bounds at all: for 510(k) clearances the 90th percentile is 10.38 months and expected shortfall at the 90 percent level 15.25 months, for original premarket approvals 27.04 and 46.36 months. Whether the surviving term is long or short in absolute years depends in addition on the date from which the term-bearing patent runs, which this article does not establish.

Identification turns the design of duration reporting into a policy variable. A median, as Implementing Regulation (EU) 2026/977 requires, can be recovered from coarse brackets but carries no information about the tail; a 90th percentile is identified within one bracket width only where the cumulative share crosses it inside a closed bracket, which the twentieth-survey counts do not do; and expected shortfall requires a closed upper boundary. An additional reporting boundary at 36 months that contained all longer durations would turn an unbounded quantity into one identified within 12 months. The same measures price a restoration instrument indexed to the tail. An extension equal to the delay beyond the 90th percentile costs in expectation $(1 - \alpha)(\text{ES}_\alpha - q_\alpha)$ of additional protection per dossier, 0.49 months for 510(k) clearances and 1.93 months for original premarket approvals in the FDA records, concentrated on the one dossier in ten that needs it; for European certification the same expression has a lower bound of zero and no upper bound, because the underlying quantile is not bounded above. Lenders and investors face the same arithmetic. Expected term shortfall is the quantity that a covenant on patent-backed financing would need, and Proposition 3 implies that a covenant calibrated on pooled statistics understates the risk carried by SME borrowers throughout the upper tail.

5 DISCUSSION AND POLICY IMPLICATIONS

Read together, the findings reposition how regulatory friction enters an account of innovation. Effective term satisfies $E[T_j] = T^{\max} - D - \mu_j$, so once mean clearance is fixed variance leaves no trace at that level of description, and whether dispersion matters depends on the frictions a firm faces. A diversified investor without financing constraints would hardly price it. An undiversified owner, a firm whose financing must bridge the wait, or a firm facing a fixed commercial window values the same distribution through a concave function, so that $E[V(T_j)] < V(E[T_j])$ and the gap widens with $\text{Var}[T_j]$. Two regimes can share a diagnosis of mean delay while the firms that matter most for European device innovation value them very differently.

Recent work that treats regulation and technology as developing together reads procedural uncertainty as largely informational and proposes infrastructures that allow regulator and firm to move from one-off review towards continuous exchange (Chin et al., 2026). The friction isolated here is capacity-bound instead. A finite pool of notified bodies processes heterogeneous device classes through shared capacity, so service-time variability propagates into the spread of waiting times, and no audit trail removes it. Better verification might compress the informational part of mean clearance time and still leave the variance across bodies and device classes untouched.

A closer structural precedent comes from political economy. Carpenter (2004) represents product approval as an optimal-stopping problem faced by a learning regulator and shows that a facially neutral regime favours larger and older firms even when the regulator is neither captured nor risk-averse, because the regulator's uncertainty about a firm falls as it accumulates submissions. The experience gradient estimated in Table 5 is consistent with that channel. Dispersion in the present account has a second source that does not contract with familiarity, a congested queue that a small firm meets on its tenth filing as on its first, and it acts after access, on the patent term that survives the wait; the SME incidence of Proposition 2 is a property of pooling rather than of behaviour. Addressing the familiarity advantage would therefore leave the capacity-driven component in place.

Aggregated over a heterogeneous population of applicants, the firm-level results imply a shift in the composition of innovation, with fewer device projects developed for the EU first, more US-first filings, and more reliance on protection mechanisms other than patents where filing is expensive and timelines are long (Ayerbe et al., 2024), and investors must price a commercialisation date that is widely distributed at the moment of investment (Bonnin Roca and O'Sullivan, 2022). Such statements concern incentives across the population and do not assert a causal effect for any single firm.

5.1 Policy implications

Article 121 MDR requires the Commission to evaluate the regulation by 27 May 2027 (European Parliament and Council, 2017). In December 2025 the Commission proposed a targeted simplification of the device rules, now in the ordinary legislative procedure as 2025/0404(COD), and the accompanying staff working document estimates annual cost savings

near EUR 3.3 billion (European Commission, 2025a,b). The proposal names the lack of sufficiently predictable certification timelines as a problem, which is independent institutional recognition of the quantity formalised here.

Implementing Regulation (EU) 2026/977, adopted on 4 May 2026, goes further in procedural terms (European Commission, 2026). Notified bodies must complete application review and contract signature within 30 days, quality-management audits within 120 days and product verification within 90 days, subject to documented interruptions; under Article 4(2) each body reports the share of activities completed within those timelines, the median duration from application to certification and the median total cost, and publishes the figures annually. The regulation applies from 25 February 2027, its reporting provision from 1 January 2028, and the monitoring covers only agreements signed after 25 May 2027.

Measured against the present analysis, the template records the centre of each body's duration distribution and a compliance share, but no measure of spread. Adding a dispersion statistic to Article 4(2), for example the 90th percentile of duration by device class, would require an amendment of an implementing act and not of the MDR itself. Even without it, the published medians of some fifty bodies will reveal the dispersion between bodies from the 2028 reporting cycle onwards. A reform that lowers median durations while the upper tail persists would, on this account, provide relief without correction.

Reporting of that kind bears on the wider question of how far European regulation can be made an instrument of competitiveness rather than a constraint upon it (Mariotti, 2024). Three design choices make a dispersion statistic workable, and Online Resource 1, Table H.2 sets what the present template records against what the measures of Section 4.2 require. Durations should be reported by device class and route, since pooling quality-management certificates for Class IIa devices with product assessments for Class III devices would reproduce at body level the aggregation problem of Proposition 2. Interruptions for missing documentation should be reported separately from the time the body consumes, so that the upper tail can be attributed to applicants or to assessors. And coverage should extend to procedures that have not yet closed, because a report restricted to completed certifications truncates exactly the long durations that matter. The identification results rank these choices. A median can be recovered from coarse brackets, a 90th percentile only within one bracket width, and expected shortfall only if reporting closes the upper tail.

Three kinds of instrument address the dispersion channel directly. Performance governance for notified bodies along the lines of the MDUFA model, with negotiated targets, public reporting and consequences for missed targets, would supply the accountability that survey monitoring measures but does not enforce (Kattel and Mazzucato, 2018; Bonnin Roca, 2024). Support differentiated for SMEs, through pre-submission advice that brings dossiers to completeness, fee relief for first-time applicants or designated review tracks, acts on the tail that the experience gradient of Table 5 documents, consistent with the view that policy for SMEs should respond to differences among them (Juergensen et al., 2020). A device-specific instrument restoring term lost to clearance

is a third option; its welfare arithmetic and legal design are developed elsewhere (Peters, 2026a, forthcoming), and Online Resource 1, Section E sketches how an extension could respond to upper-tail delay.

Any argument that clearance time erodes appropriability invites the rejoinder that the erosion is the price of safety. The relevant counterfactual is the predecessor regime with its own record of harm, not a frictionless world, and the asymmetry of the trade-off bears on policy. The appropriability loss is measurable, immediate and concentrated on the firms identified by P2 and Table 5; the safety benefit set against it remains hard to demonstrate (Bretthauer et al., 2023; McDermott and Kearney, 2025), accrues over the deployed life of devices and is spread across the patient population. A cost with that profile, weighed against a benefit of the opposite profile, does not argue for dismantling scrutiny; it argues for delivering the same scrutiny with less dispersion. Welfare accounting has to include access as well, since changes in exclusivity alter the availability of health technologies (Bennato and Giulietti, 2019; Peters, 2026a), and longer assessment can generate information about quality (Grennan and Town, 2020), so not every reduction of dispersion raises welfare. The argument concerns dispersion that stems from class mix, resubmission and congestion, which delays access without adding information about the device.

Asymmetry in the trade-off bears directly on policy. On the cost side, the appropriability loss is measurable in the erosion and gap bands of R1, immediate because it falls on devices now in the queue, and concentrated on the SMEs identified by P2. The safety benefit set against it is harder to demonstrate, with post-MDR evidence on improved outcomes still thin (Bretthauer et al., 2023; McDermott and Kearney, 2025); it accrues over the deployed life of devices; and it is spread across the patient population. A cost with that profile, weighed against a benefit of the opposite profile, does not argue for dismantling scrutiny. It argues for delivering the same scrutiny with less variance. Welfare accounting has to include the availability side as well, since changes in exclusivity alter access to health technologies (Bennato and Giulietti, 2019; Peters, 2026a).

Firm size is the incidence dimension formalised here, but not the only one, since geography and product category cut across the same population (Dobrzynska et al., 2025; Gupta et al., 2025), and the conditions under which the construct travels to other regulated sectors are set out in Online Resource 1, Section E.

5.2 Boundaries and limitations

Several limitations bound the results. The queueing relation serves as a comparative-statics argument with a single server standing in for some fifty bodies, and a multi-server representation is not modelled. EU durations are available only as shares of notified bodies by typical duration and as firm-reported phase durations from small samples, so file-level dispersion in the EU is not identified and log-normal and heavier-tailed families cannot be told apart at the upper tail. Submission frequency in the US record classifies firms by how often they appear in the register and measures neither experience at the moment of filing nor firm size. A strictly backward-looking count reproduces the gradient; a within-

applicant specification on receipt cohorts does not, so what the estimates identify is a difference between firms and not learning inside one, and the associations carry no causal reading. Absolute effective terms are not reported, because the date from which a term-bearing patent runs is the earliest application in its chain of benefit and establishing that date for a sample of devices is a task in patent-register work not undertaken here; revenue is proxied by turnover per enterprise, which includes exports and other products; financing intervals and dilution come from venture data across sectors; and the cost of a financing round is measured as dilution times sector valuation multiples, not as an observed loss. Strategic interaction between firms and between notified bodies is abstracted, EU and US clearance times are treated as independent (positive correlation would strengthen the case for US-first entry), and quality benefits of the MDR are not incorporated. The value of waiting for information and spillovers from US clinical data are not computed. Proposition 3 requires upper-tail dominance of SME clearance times, which the firm-size data support above about ten months of review, and the bounds of Online Resource 1, Table H.3 depend on the cap assumed for the open upper bracket.

Each limitation marks a test. Matched firm-level records linking patent filing dates, FDA submissions and MDR applications would allow the sequencing condition to be confronted with observed entry decisions. Annual notified-body reports under Implementing Regulation (EU) 2026/977 will provide body-level medians and compliance shares from 2028, and processing times cross-tabulated by firm size would allow the mixture identity to be estimated from administrative records. A further line concerns devices that re-enter assessment after material changes, as adaptive software increasingly does; there the single clearance draw becomes a repeated one and the variance channel recurs over a device's life, a case reserved for separate work.

Pharmaceutical patents face structurally similar delay problems, but statutory restoration through patent-term extension in the United States and supplementary protection certificates in Europe, whose compatibility with the European patent system Article 63(2)(b) EPC provides for, insulates them partly (Grabowski et al., 2015; European Parliament and Council, 2009). Standalone devices lack an equivalent mechanism and bear the full appropriability cost.

6 CONCLUDING REMARKS

Administrative time is a price that firms pay in a currency regulators rarely measure. Where the procedure at issue is a condition of market access, that price is levied on a wasting asset: patent term runs from filing, so a month in conformity assessment is a month of exclusivity that no later efficiency can restore. Firms differ in what such a month costs them. A diversified owner with cash on hand discounts it and moves on; a single-product manufacturer financing the wait from a fixed runway may find that the tail of the distribution, and not its centre, decides whether the project survives at all. Such an asymmetry gives the dispersion of regulatory duration an economic content distinct from its level, and it gives reporting rules a role they are not usually credited with, since a statistic that describes only the centre cannot show a burden that lives

in the tail. For the European medical-device sector, where small and medium-sized manufacturers hold a minority of devices but require a plurality of certificates, the question is not academic.

Against that background the author has measured what the clearance clock does to appropriable patent term under the MDR and how the effect is distributed across firm types. The measurement rests on four public records, namely the notified-body survey commissioned by the European Commission, with data status 28 February 2026; the manufacturer surveys of MedTech Europe for 2022 and 2024, which report durations and fees separately by firm size; the file-level databases of the US Food and Drug Administration, covering 15,706 premarket notifications, 173 De Novo grants and 311 original premarket approvals; seventy Federal Register notices determining regulatory review periods; and the patent office's list of extensions granted under 35 U.S.C. §156, which records what 111 of the 223 patents named in all 161 device notices actually recovered. Costs of capital, valuation multiples, financing intervals and turnover by size class complete the inputs.

Three results follow. European certification consumes about one to 1.4 years more patent term than a 510(k) clearance and 0.3 to 0.6 years more than an original premarket approval, and it roughly doubled relative to the predecessor directives; a device that passes premarket approval is in addition eligible for statutory term restoration in the United States, which the patent office granted to 111 device patents at a median of 2.62 years of the review period, and no European counterpart exists. The US records identify the tails of clearance time exactly, whereas the European record bounds them only from below, so whether the European clock is also more dispersed cannot be settled from public data. Entering the United States first is never optimal for a firm that would file the same European dossier anyway, and with observed inputs it pays behind a 510(k) review only where completing that dossier saves between two and five and a half months, an amount of the order documented between smaller and larger firms, and rarely pays behind a premarket approval. Smaller manufacturers wait hardly longer on average, but their review phase has a far longer tail, 34.3 against 19.5 months at the ninetieth percentile, and the fees weigh about fourteen times more heavily against turnover than for firms with 50 to 249 persons employed; pooled statistics conceal the tail at every level above the point where the two firm-size distributions cross. On the file-level US record the same pattern is estimated rather than read from a survey: applicants with a single premarket notification in five years wait 34 percent longer than the most frequent applicants at the median and 42 percent longer at the ninth decile, with device area, submission type and fiscal year held fixed. The gradient survives a strictly backward-looking experience count but not a within-applicant specification on receipt cohorts, so it separates firms rather than tracking learning inside one. Read together, the results locate the burden where pooled monitoring cannot see it.

Regulators should measure dispersion, not only duration. Implementing Regulation (EU) 2026/977 obliges notified bodies to report the share of activities completed within the new maximum timelines and the median duration from application to certification; neither statistic describes the tail that

decides whether a small manufacturer can finance the wait. Adding a ninetieth percentile by device class and conformity-assessment route, and closing the upper reporting bracket, would make the quantity observable at negligible cost and would give the evaluation required by Article 121 MDR a criterion that a reform can be judged against. Where the tail is what binds, the tail is what must be reported.

DECLARATIONS

Funding. No funding was received for conducting this study.

Competing interests. The author declares no competing financial interests. The analysis draws exclusively on publicly available regulatory and industry data; no proprietary or client-specific information shaped it.

Ethics approval and consent. Not applicable. The study uses publicly available aggregate and file-level administrative data and involves no human participants.

Data and code availability. All data are public. Online Resource 2 contains the file-level FDA extracts, the parsed Federal Register notices, the granted extensions, the notified-body survey series, the bracket definitions, the receipt-cohort records with the reconstructed applicant crosswalk, the Eurostat size-class figures, a table of every externally sourced parameter with its source, and a Python script that regenerates every computed quantity reported here. The same package is deposited at figshare under <https://doi.org/10.6084/m9.figshare.33683083>, under a CC BY 4.0 licence. The archive there has the SHA-256 digest a9a853495b620b3ee67e3a16d8eb6888f40f17678a193ae0f4fe7377e63a497c and carries a manifest of the seventeen files inside it, so that a reader can establish both that the download is intact and that the copy being run is the one the reported quantities come from.

Author contribution. The author is solely responsible for the conception, the analysis and the writing.

Supplementary information. Online Resource 1 (PDF) contains the derivations and proofs, the estimation detail, the robustness results, the supplementary tables and figures and the glossary. Online Resource 2 (ZIP) is the reproduction package described above.

REFERENCES

- Acerbi, C., & Tasche, D. (2002). On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7), 1487–1503. [https://doi.org/10.1016/S0378-4266\(02\)00283-2](https://doi.org/10.1016/S0378-4266(02)00283-2)
- Agiropoulos, C., Drivas, K., Galanos, G., & Agiakloglou, C. (2026). How cannibalization effects shape trademark renewal decisions at the firm level. *Journal of Industrial and Business Economics*. <https://doi.org/10.1007/s40812-026-00427-x>
- Antonelli, C. (2017). Digital knowledge generation and the appropriability trade-off. *Telecommunications Policy*, 41(10), 991–1002. <https://doi.org/10.1016/j.telpol.2016.12.002>
- Artzner, P., Delbaen, F., Eber, J.-M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228. <https://doi.org/10.1111/1467-9965.00068>
- Asimakopoulou, G., & Paraskevopoulou, E. (2025). Incumbents' adaptation in institutionally induced transitions: entry timing, institutional experience and performance. *Journal of Industrial and Business Economics*, 53(1), 143–166. <https://doi.org/10.1007/s40812-025-00353-4>

- Austrian National Public Health Institute, Areté, & Civic Consulting (2026a). *Study supporting the monitoring of the availability of medical devices on the EU market: survey results of the 18th NB survey (MDR/IVDR) with data status 31 October 2025*. Commissioned by the European Commission within the EU4Health Programme, 25 March 2026.
- Austrian National Public Health Institute, Areté, & Civic Consulting (2026b). *PowerPoint presentation containing a study overview and survey results of the 20th NB survey for the “Study supporting the monitoring of availability of medical devices on the EU market”* (data status 28 February 2026). Commissioned by the European Commission within the EU4Health Programme, 2 July 2026. https://health.ec.europa.eu/medical-devices-topics-interest/notified-bodies-medical-devices_en
- Aydin Ozden, S., & Khashabi, P. (2023). Patent remedies and technology licensing: evidence from a Supreme Court decision. *Strategic Management Journal*, 44(9), 2311–2338. <https://doi.org/10.1002/smj.3493>
- Ayerbe, C., Boulos, C., & Castellaneta, F. (2024). Navigating protection mechanisms and innovation models: a literature-based configurational framework of intellectual property strategies. *Technovation*, 137, 103101. <https://doi.org/10.1016/j.technovation.2024.103101>
- Battaggion, M. R., Garrod, L., Giulietti, M., & Grilli, L. (2023). Emerging challenges in competition policy and regulation. *Journal of Industrial and Business Economics*, 50(2), 221–225. <https://doi.org/10.1007/s40812-023-00269-x>
- Benassi, M., Grinza, E., & Rentocchini, F. (2020). The rush for patents in the Fourth Industrial Revolution. *Journal of Industrial and Business Economics*, 47(4), 559–588. <https://doi.org/10.1007/s40812-020-00159-6>
- Bennato, A. R., & Giulietti, M. (2019). Patent policy regulation and public health. *Journal of Industrial and Business Economics*, 46(4), 431–457. <https://doi.org/10.1007/s40812-019-00131-z>
- Boakye, D., Sarpong, D., & Mordi, C. (2022). Regulatory review of new product innovation: conceptual clarity and future research directions. *Technological Forecasting and Social Change*, 175, 121419. <https://doi.org/10.1016/j.techfore.2021.121419>
- Bonnin Roca, J. (2024). Regulatory agencies as innovation enablers: a conceptualization. *Science and Public Policy*, 51(6), 1051–1061. <https://doi.org/10.1093/scipol/scae049>
- Bonnin Roca, J., & O’Sullivan, E. (2022). The role of regulators in mitigating uncertainty within the Valley of Death. *Technovation*, 109, 102157. <https://doi.org/10.1016/j.technovation.2020.102157>
- Brethauer, M., Gerke, S., Hassan, C., Ahmad, O. F., & Mori, Y. (2023). The new European Medical Device Regulation: balancing innovation and patient safety. *Annals of Internal Medicine*, 176(6), 844–848. <https://doi.org/10.7326/M23-0454>
- Budish, E., Roin, B. N., & Williams, H. (2015). Do firms underinvest in long-term research? Evidence from cancer clinical trials. *American Economic Review*, 105(7), 2044–2085. <https://doi.org/10.1257/aer.20131176>
- Calvino, F., Criscuolo, C., & Verlhac, R. (2020). *Declining business dynamism: structural and policy determinants*. OECD Science, Technology and Industry Policy Papers No. 94. Paris: OECD Publishing. <https://doi.org/10.1787/77b92072-en>
- Capuano, C., & Grassi, I. (2019). Imperfect patent protection, licensing, and willingness to pay for the innovation. *Journal of Industrial and Business Economics*, 47(2), 333–359. <https://doi.org/10.1007/s40812-019-00142-w>
- Carl, A.-K., & Hochmann, D. (2024). Impact of the new European medical device regulation: a two-year comparison. *Biomedical Engineering / Biomedizinische Technik*, 69(3), 317–326. <https://doi.org/10.1515/bmt-2023-0325>
- Carpenter, D. P. (2004). Protection without capture: product approval by a politically responsive, learning regulator. *American Political Science Review*, 98(4), 613–631. <https://doi.org/10.1017/S0003055404041383>
- Carta (2025). *State of private markets: Q2 2025*. <https://carta.com/data/state-of-private-markets-q2-2025/>
- Carta (2026). *State of private markets: 2025 in review*. <https://carta.com/data/state-of-private-markets-q4-2025/>
- Chin, T., Li, Z., & Caputo, F. (2026). From regulatory review to co-evolution: blockchain as regulatory infrastructure for business model innovation. *Technological Forecasting and Social Change*, 229, 124727. <https://doi.org/10.1016/j.techfore.2026.124727>
- Ciriaci, D. (2014). *Business dynamics and red tape barriers*. European Economy Economic Papers No. 532. Brussels: European Commission, Directorate-General for Economic and Financial Affairs. <https://doi.org/10.2765/71352>
- Corsino, M., & Torrisi, S. (2023). University engagement in open innovation and intellectual property: evidence from university–industry collaborations. *Journal of Industrial and Business Economics*, 50(4), 781–813. <https://doi.org/10.1007/s40812-023-00280-2>
- Court of Justice of the European Union (2018). Judgment of 25 October 2018, *Boston Scientific Ltd v Deutsches Patent- und Markenamt*, Case C-527/17, EU:C:2018:867.
- Czarnitzki, D., & Toole, A. A. (2011). Patent protection, market uncertainty, and R&D investment. *Review of Economics and Statistics*, 93(1), 147–159. https://doi.org/10.1162/REST_a_00069
- Daigle, B., & Torsenaar, M. (2019). *The EU Medical Device Regulation and the U.S. medical device industry*. Executive Briefings on Trade. Washington, DC: U.S. International Trade Commission.
- Damodaran, A. (2026). *Cost of equity and capital, and price and value to sales ratios, United States and Europe, data updated January 2026*. New York University Stern School of Business. <https://pages.stern.nyu.edu/>

- ~adamodar/New_Home_Page/data.html
- Davis, M. H. A. (1993). *Markov models and optimization*. London: Chapman & Hall. <https://doi.org/10.1007/978-1-4899-4483-2>
- Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton: Princeton University Press.
- Djankov, S., McLiesh, C., & Ramalho, R. M. (2006). Regulation and growth. *Economics Letters*, 92(3), 395–401. <https://doi.org/10.1016/j.econlet.2006.03.021>
- Dobrzynska, A., Rejón-Parrilla, J. C., Epstein, D., Aranda-López, J., Fraser, A. G., & Blasco-Amaro, J. A. (2025). Survey of notified bodies reveals very limited use of conditional certification for high-risk medical devices. *Frontiers in Medical Technology*, 7, 1504294. <https://doi.org/10.3389/fmedt.2025.1504294>
- European Commission (2025a). Proposal for a Regulation of the European Parliament and of the Council amending Regulations (EU) 2017/745 and (EU) 2017/746 as regards simplifying and reducing the burden of the rules on medical devices and in vitro diagnostic medical devices. COM(2025) 1023 final, 16 December 2025. Procedure 2025/0404(COD).
- European Commission (2025b). Commission Staff Working Document accompanying the proposal COM(2025) 1023 final. SWD(2025) 1050 final, 16 December 2025.
- European Commission (2026). Commission Implementing Regulation (EU) 2026/977 of 4 May 2026 laying down certain uniform quality management and procedural requirements for the conformity assessment activities carried out by a notified body designated under Regulations (EU) 2017/745 and (EU) 2017/746. *Official Journal of the European Union*, L, 2026/977, 5 May 2026.
- European Parliament and Council (2009). Regulation (EC) No 469/2009 of 6 May 2009 concerning the supplementary protection certificate for medicinal products. *Official Journal of the European Union*, L 152, 1–10.
- European Parliament and Council (2017). Regulation (EU) 2017/745 of 5 April 2017 on medical devices. *Official Journal of the European Union*, L 117, 1–175.
- Eurostat (2026). *Enterprise statistics by size class and NACE Rev. 2 activity (from 2021 onwards)*, dataset sbs_sc_ovw, update of 1 September 2026. https://ec.europa.eu/eurostat/api/dissemination/statistics/1.0/data/sbs_sc_ovw
- Federal Register (2026). *Determination of regulatory review period for purposes of patent extension*, notices of the Food and Drug Administration for medical devices published 9 February 2018 to 26 February 2026, retrieved through the Federal Register API. <https://www.federalregister.gov/api/v1/documents.json>
- Fleming, W. H., & Soner, H. M. (2006). *Controlled Markov processes and viscosity solutions* (2nd ed.). New York: Springer. <https://doi.org/10.1007/0-387-31071-1>
- Food and Drug Administration (2026a). *Premarket notification 510(k) database*, file pmn96cur.zip, accessed 11 September 2026. <https://www.accessdata.fda.gov/premarket/ftparea/pmn96cur.zip>
- Food and Drug Administration (2026b). *Premarket approval database*, file pma.zip, accessed 11 September 2026. <https://www.accessdata.fda.gov/premarket/ftparea/pma.zip>
- Grabowski, H. G., DiMasi, J. A., & Long, G. (2015). The roles of patents and research and development incentives in biopharmaceutical innovation. *Health Affairs*, 34(2), 302–310. <https://doi.org/10.1377/hlthaff.2014.1047>
- Granstrand, O. (2024). A cooperative game theory systems approach to the value analysis of (innovation) ecosystems. *Technovation*, 130, 102926. <https://doi.org/10.1016/j.technovation.2023.102926>
- Granstrand, O., & Holgersson, M. (2020). Innovation ecosystems: a conceptual review and a new definition. *Technovation*, 90–91, 102098. <https://doi.org/10.1016/j.technovation.2019.102098>
- Grennan, M., & Town, R. J. (2020). Regulating innovation with uncertain quality: information, risk, and access in medical devices. *American Economic Review*, 110(1), 120–161. <https://doi.org/10.1257/aer.20180946>
- Grossman, G. M., & Lai, E. L.-C. (2004). International protection of intellectual property. *American Economic Review*, 94(5), 1635–1653.
- Guoru, T., Keshminder, J. S., Zaghol, A. K., & Ridzuan, A. R. (2026). The impact of economic policy uncertainty on innovation in Chinese enterprises. *Journal of Industrial and Business Economics*, 53(3), 905–940. <https://doi.org/10.1007/s40812-026-00390-7>
- Gupte, A., Nitave, R., & Gobburu, J. (2025). Regulatory landscape of accelerated approval pathways for medical devices in the United States and the European Union. *Frontiers in Medical Technology*, 7, 1586070. <https://doi.org/10.3389/fmedt.2025.1586070>
- Holgersson, M. (2013). Patent management in entrepreneurial SMEs: a literature review and an empirical study of innovation appropriation, patent propensity, and motives. *R&D Management*, 43(1), 21–36. <https://doi.org/10.1111/j.1467-9310.2012.00700.x>
- Holgersson, M., Granstrand, O., & Bogers, M. (2018). The evolution of intellectual property strategy in innovation ecosystems: uncovering complementary and substitute appropriability regimes. *Long Range Planning*, 51(2), 303–319. <https://doi.org/10.1016/j.lrp.2017.08.007>
- Huuskio, J., Kinnunen, U.-M., & Saranto, K. (2023). Medical device regulation (MDR) in health technology enterprises: perspectives of managers and regulatory professionals. *BMC Health Services Research*, 23, 310. <https://doi.org/10.1186/s12913-023-09316-8>
- Johnson, G. A., Shriver, S. K., & Goldberg, S. G. (2023). Privacy and market concentration: intended and unintended consequences of the GDPR. *Management Science*, 69(10), 5695–5721. <https://doi.org/10.1287/mnsc.2023.4709>
- Juergensen, J., Guimón, J., & Narula, R. (2020). European

- SMEs amidst the COVID-19 crisis: assessing impact and policy responses. *Journal of Industrial and Business Economics*, 47(3), 499–510. <https://doi.org/10.1007/s40812-020-00169-4>
- Kattel, R., & Mazzucato, M. (2018). Mission-oriented innovation policy and dynamic capabilities in the public sector. *Industrial and Corporate Change*, 27(5), 787–801. <https://doi.org/10.1093/icc/dty032>
- Kleinrock, L. (1975). *Queueing systems, Vol. I: Theory*. New York: Wiley.
- Malerba, F. (2002). Sectoral systems of innovation and production. *Research Policy*, 31(2), 247–264. [https://doi.org/10.1016/S0048-7333\(01\)00139-1](https://doi.org/10.1016/S0048-7333(01)00139-1)
- Manski, C. F. (2003). *Partial identification of probability distributions*. New York: Springer. <https://doi.org/10.1007/b97478>
- Marco, A. C., & Prieger, J. E. (2009). Congestion pricing for patent applications. SSRN Working Paper 1443470. <https://doi.org/10.2139/ssrn.1443470>
- Mariotti, S. (2024). “Open strategic autonomy” as an industrial policy compass for the EU competitiveness and growth: the good, the bad, or the ugly? *Journal of Industrial and Business Economics*, 52(1), 1–26. <https://doi.org/10.1007/s40812-024-00327-y>
- Mazzucato, M. (2018). Mission-oriented innovation policies: challenges and opportunities. *Industrial and Corporate Change*, 27(5), 803–815. <https://doi.org/10.1093/icc/dty034>
- McDermott, O., & Kearney, B. (2025). A review of the literature on the new European Medical Device Regulations requirements for increased clinical evaluation. *International Journal of Pharmaceutical and Healthcare Marketing*, 19(1), 1–21. <https://doi.org/10.1108/IJPHM-07-2023-0060>
- MedTech Europe (2022). *Survey report: availability of medical devices in connection to MDR implementation*. Brussels: MedTech Europe.
- MedTech Europe (2024). *IVDR & MDR survey results 2024: public report*. Brussels: MedTech Europe. https://www.medtecheurope.org/wp-content/uploads/2025/01/mte_report_ivdr_mdr_2024-v7.pdf
- MedTech Europe (2025). *IVDR & MDR survey results 2024: highlights report*. Brussels: MedTech Europe, January 2025.
- Peltzman, S. (1973). An evaluation of consumer protection legislation: the 1962 drug amendments. *Journal of Political Economy*, 81(5), 1049–1091. <https://doi.org/10.1086/260107>
- Peters, A. (2026a). Regulatory delays and patent term erosion in medical device innovation: a two-market analysis of EU and US frameworks. *European Journal of Law and Economics*, 62(1), 131–164. <https://doi.org/10.1007/s10657-026-09873-3>
- Peters, A. (2026b). Effective patent time left: a reproducible legal-status metric for U.S. medical devices. *World Patent Information*, 87, 102487. <https://doi.org/10.1016/j.wpi.2026.102487>
- Peters, A. (forthcoming). A device protection certificate for medical devices under EU law: doctrinal feasibility and legislative design after Boston Scientific. *GRUR International*.
- Somaya, D. (2012). Patent strategy and management: an integrative review and research agenda. *Journal of Management*, 38(4), 1084–1114. <https://doi.org/10.1177/0149206312444447>
- Stern, A. D. (2017). Innovation under regulatory uncertainty: evidence from medical technology. *Journal of Public Economics*, 145, 181–200. <https://doi.org/10.1016/j.jpubeco.2016.11.010>
- Svempe, L. (2024). Exploring impediments imposed by the Medical Device Regulation EU 2017/745 on software as a medical device. *JMIR Medical Informatics*, 12, e58080. <https://doi.org/10.2196/58080>
- Teece, D. J. (1986). Profiting from technological innovation: implications for integration, collaboration, licensing and public policy. *Research Policy*, 15(6), 285–305. [https://doi.org/10.1016/0048-7333\(86\)90027-2](https://doi.org/10.1016/0048-7333(86)90027-2)
- Tomasi, C., Pieri, F., & Cecco, V. (2023). Red tape and industry dynamics: a cross-country analysis. *Journal of Industrial and Business Economics*, 50(2), 283–320. <https://doi.org/10.1007/s40812-023-00266-0>
- United States Patent and Trademark Office (2026). *Patent terms extended under 35 U.S.C. §156*, file pte_certs.xls, accessed 12 September 2026. <https://www.uspto.gov/patents/laws/patent-term-extension/patent-terms-extended-under-35-usc-156>
- US Food and Drug Administration (2022). *MDUFA performance goals and procedures, fiscal years 2023 through 2027* (commitment letter). Silver Spring, MD: FDA. <https://www.fda.gov/media/158308/download>
- US Food and Drug Administration (2025). *Performance report to Congress: MDUFA, FY 2024*. Silver Spring, MD: FDA.
- US Food and Drug Administration (2026). *Performance report to Congress: MDUFA, FY 2025* (data as of 30 September 2025). Silver Spring, MD: FDA. <https://www.fda.gov/media/191127/download>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [ESM1.pdf](#)